

AFIT/GE/ENG/95D-17

Detection of Clustered Microcalcifications Using Wavelets

THESIS


Donald A. McCandless
1st Lieutenant, USAF

AFIT/GE/ENG/95D-17

19960327 042

Approved for public release; distribution unlimited

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

AFIT/GE/ENG/95D-17

Detection of Clustered Microcalcifications Using Wavelets

THESIS

Presented to the Faculty of the School of Engineering
of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Electrical Engineering

Donald A. McCandless, B.S. Electrical Engineering
1st Lieutenant, USAF

December 1995

Approved for public release; distribution unlimited

Acknowledgements

No man is an island, and this thesis would not have been possible without the help and encouragement of several individuals over the course of the past year. First I owe a great deal to my advisor Dr. Rogers for his enthusiasm and patient encouragement, and for helping me catch the vision of this research. I also wish to thank my committee members, Maj. Ruck and Capt. Raines, for their inputs and taking the time to make themselves available to provide any assistance. Special thanks to Maj. Jeffrey Hoffmeister, who served as the medical expert, and without his knowledge this whole undertaking would have been unimaginable. Thanks to my partners-in-crime, Bill, Ron, Amy, and Dave for their help and advice, and especially in the effort to acquire enough images to make this research possible. I am also grateful for the network support provided so ably by Dan Zambon and Dave Doak, whose patience with our disk storage requirements has been saintly. And finally, I would like to thank my wife Sheri, whose sacrifice and caring support over the past 18 months has uplifted me and made me a better person.

Donald A. McCandless

Table of Contents

	Page
Acknowledgements	ii
List of Figures	vi
List of Tables	viii
Abstract	ix
I. Introduction	1
1.1 Cancer Facts	1
1.2 Computer-aided Diagnosis	1
1.3 Problem Statement	2
1.4 Scope	2
1.5 Approach	2
1.6 Objectives	3
1.7 Organization	3
II. Background	4
2.1 Introduction	4
2.2 Multiresolution/Wavelets	4
2.3 Feature Extraction	6
2.4 Model-Based Vision	8
2.5 Conclusion	8
III. Methods	10
3.1 Introduction	10
3.2 Database	10
3.3 Method Overview	10

	Page
3.4 Focus of Attention	12
3.4.1 Gray Level Modification	12
3.4.2 Filtering	15
3.4.3 Region of Interest Identification	18
3.5 Indexing	19
3.5.1 Index Features	20
3.5.2 Index Decision-making	24
3.6 Matching	25
3.6.1 Final Object Identification	25
3.6.2 Hypothesis Testing	28
3.7 Conclusion	29
IV. Results	30
4.1 Introduction	30
4.2 Training Set	30
4.2.1 Focus of Attention	30
4.2.2 Indexing	31
4.2.3 Matching	32
4.2.4 Discussion	32
4.3 Test Set	35
4.4 Evaluation Set	35
4.5 Normal Tissue Set	36
4.6 Summary	36
V. Conclusion	38
5.1 Introduction	38
5.2 Summary of Methodology	38
5.3 Summary of Results	39
5.4 Conclusion	41

	Page
Appendix A. Wavelet Code	43
Appendix B. Microcalcification Detection Code	44
Appendix C. Feature Selection Code	45
Appendix D. Wavelet Code	46
D.1 Introduction	46
D.2 Requirements	46
D.3 new_wave.m	46
D.4 new_recon.m	48
Appendix E. Microcalcification Detection Code	50
E.1 Introduction	50
E.2 main_seg_roc.m	50
E.3 pre_proc64.m	52
E.4 hypoth.m	55
E.5 roi_correl64.m	56
E.6 reduce64.m	59
E.7 roc_feat64.m	59
E.8 sorter64.m	61
Appendix F. Feature Selection Code	63
F.1 Introduction	63
F.2 classified.m	63
F.3 parzen_prob_fast.m	64
F.4 find_bound.m	64
F.5 get_edbfm.m	66
Bibliography	67
Vita	70

List of Figures

Figure	Page
1. Entire procedure for detecting microcalcification clusters, starting with the digitized mammogram and ending with a set of ROIs	11
2. (a) Original digitized mammogram. (b) Histogram of the gray levels for the mammogram in (a) - the spike near 800 corresponds to the dark background around the breast, most of the relevant information is in the range from 2200 to 3600. (c) Comparison between the linear and non-linear relationship of pixel value to gray-scale value. For the non-linear relationship, the pixels in the steepest portion of the curve have the highest dynamic range and contrast.	13
3. 64-by-64 pixel ROI containing microcalcifications. (a) Original (unmodified) pixel values (b) rescaled ROI using the non-linearity showing increased contrast	14
4. Block diagram of one level of (a) wavelet decomposition and (b) wavelet reconstruction [16].	16
5. Wavelet decomposition to (a) first level (b) third level.	17
6. Mammogram that is wavelet decomposed to third level.	17
7. Demonstration of centroid migration (a) three separate windows each contain a portion of the microcalcification cluster; (b) final ROI after several successive centroid migrations containing the entire cluster.	18
8. Cluster identification within an ROI: (a) Original filtered ROI, (b) ROI after thresholding with a high value, (c) ROI after thresholding with a lower value, (d) result after several iterations of morphological dilation of the the ROI in (b) and logical ANDing with the ROI in (c). The remaining four objects are then evaluated to determine their likelihood of being microcalcifications.	20
9. Mesh plots of (a) filtered microcalcification object (b) original microcalcification (c) auto-correlation of filtered microcalcification (d) cross-correlation of the filtered and original.	23
10. Mesh plots of (a) filtered "microcalcification" object (b) original image showing a tissue edge (c) auto-correlation of filtered microcalcification (d) cross-correlation of the filtered and original.	24

Figure	Page
11. Comparison of object separation using the original ROI: (a) & (b) without removing the non-stationary background; (c) & (d) after removing the non-stationary background. Once the non-stationary portion of the background is removed, the dimmer microcalcifications are retained, and the background object on the bottom of the original image is discarded.	26
12. Illustration of how the final ROI which is used for hypothesis testing is generated from the two filtered versions of the ROI.	27
13. (a) Decision boundary and data points: * = phantom data points, + = actual microcalcifications, o's = noise data points. (b) Blow-up of the critical portion of the decision boundary and data points actually used to determine the decision boundary: + and o indicate data points, x's are decision boundary points as determined by Lee and Landgrebe's non-parametric procedure	28
14. ROC curves showing performance as detection parameters are changed: (a) multiple of the standard deviation used for segmentation, (b) initial low-hi threshold pair, (c) correlation peak and mean square error criteria, (d) hi-pass ROI threshold	33
15. Composite probability of detection and false positive ROIs per image results of training, test, and evaluation data	37

List of Tables

Table		Page
1.	Table showing the contrast and dynamic range improvement produced by using a non-linear function to re-scale the gray levels. Values were obtained over a sample space of fourteen images, see Equations 2 and 3 for definitions of CII and DRII.	15
2.	Feature decision thresholds and their order of consideration.	25
3.	Table showing the results of the focus of attention procedure. Num. regions indicates the number of ROIs that survived the thresholding procedure, and rank indicates the rank order by energy of the true ROI (1 = highest energy, etc.).	31
4.	Table showing the results of the indexing procedure. Mean and standard deviation numbers in parentheses are calculated by ignoring image 008.	32
5.	Number of (false) detections that were found by the algorithm in radiologist declared normal images.	36
6.	Detection and false positive results for all data sets.	40

Abstract

Identifying microcalcification clusters plays an important role in early breast cancer detection. With the increasing number of mammograms that radiologists must examine each year, computer-aided diagnosis may provide a vital second opinion to radiologists by helping them screen mammograms for early signs of cancer. An automated method for detecting microcalcification clusters is presented. The algorithm begins with a digitized mammogram and outputs the center coordinates of regions of interest (ROIs) that contain microcalcification clusters. The method presented uses a 12-tap Least Asymmetric Daubechies (LAD12) wavelet in a tree structured filter bank to increase the signal to noise level of microcalcifications. The signal to noise level gain achieved by the filtering allows subsequent thresholding to eliminate on average 90% of the image from further consideration without eliminating actual microcalcifications 95% of the time. A novel approach to isolating individual calcifications from background tissue through non-stationary noise reduction, low/hi thresholding, and morphological filtering is demonstrated; this technique reduces the number of false detections by an average of 5 per image. Several features are extracted from each potential calcification, including two newly proposed correlation features, to distinguish actual microcalcifications from correlated background tissue. Altogether, the method successfully detected 44 of 53 microcalcification clusters (83%) with an average of 2.3 false positive clusters per image. A cluster is considered detected if it contains 3 or more microcalcifications within a 6.4 mm by 6.4 mm area. Although the emphasis is placed on detecting microcalcification clusters for further examination by a radiologist with no attempt made to diagnose the cluster as malignant or benign, the method successfully detected 13 of 14 (93%) malignant cases.

Detection of Clustered Microcalcifications Using Wavelets

I. Introduction

1.1 Cancer Facts

Breast cancer is a topic of enormous concern. The National Cancer Institute (NCI) estimated that 182,000 women would be newly diagnosed with breast cancer in 1994, with approximately 46,000 deaths from the disease[1]. Early detection of breast cancer is vital to reducing mortality; five year survival rates are 93 percent for detection of localized cancer, 72 percent for regional, and only 18 percent for distant or metastatic disease[32]. Screen/film mammography has been used for many years by radiologists to screen for breast cancer. Although mammography is currently the best method for breast cancer screening, 10 percent of breast cancers, even lumps that can be felt, do not show up on x-rays[1]. In addition, 10 to 30 percent of negative readings are later proven to be from patients with breast cancer, *two-thirds of which are evident in the mammogram in retrospect*[12] (italics added). Possible explanations of why mammography fails to detect so many cancers can be attributed to several factors, including poor image quality, radiologist fatigue, and human oversight. To counter some of these factors it has been suggested that a second opinion may be the answer[12].

1.2 Computer-aided Diagnosis

It has been proposed that computer-aided diagnosis (CADx) techniques may be used to increase the speed and accuracy of breast cancer screening. Computer-aided diagnosis can be defined as a diagnosis made with input from computer analysis of radiographic features. Recent developments in computer and digital image technology have allowed for the digital processing of mammograms. These recent developments make it possible to use sophisticated computer algorithms to identify potentially cancerous regions within the breast. It is not the purpose of a CADx system to replace human radiologists, but rather to assist them.

A Model Based Vision approach to CADx will provide improved performance in a cancer detection system in terms of producing a higher probability of detection with better false positive rejection. Model based techniques use explicit models of objects, sensors, and image formation to

recognize objects in a noisy environment[2]. Focus of attention (FOA), indexing, and matching are the major steps in a model based process. Application of these steps to breast cancer detection will result in identifying and labeling potentially cancerous breast tissue.

Of particular interest to researchers is the problem of detecting tiny clusters of microcalcifications; they are generally more difficult for a radiologist to spot and are the earliest mammographically visible sign of breast cancer. Therefore, research identifying microcalcifications may reap the biggest benefit from computer aided techniques. The problem of distinguishing cancerous and pre-cancerous microcalcifications from normal tissue is complicated due to the small size of the cancerous features and the subtle differences between cancerous and non-cancerous tissue.

1.3 Problem Statement

This thesis uses multiresolution/wavelet analysis and feature extraction in a model based target recognition architecture to identify regions of interest (ROIs) within an image that contain a cluster of microcalcifications.

1.4 Scope

This work concentrates on identifying regions containing small clusters of microcalcifications; no attempt is made to identify other lesions such as masses or tumors. Also, no attempt is made to diagnose the region as malignant or benign; only to classify tissue as being normal or abnormal and in need of subsequent investigation by a radiologist. Potentially cancerous tissue is distinguished from normal tissue by analyzing features of a given region within a mammogram. Therefore, feature extraction and saliency play an important role in cancer detection; several features are presented and their classification ability is investigated.

1.5 Approach

This research presents a model based solution to the problem of detecting microcalcification clusters within breast tissue. It begins by applying a wavelet based technique to increase the signal to noise ratio (SNR) of microcalcifications and segmenting regions of interest (ROIs) that contain them. Individual objects inside the ROIs are isolated using thresholding and morphological filtering.

Features are extracted from each object and used to form a hypothesis as to whether the object is a microcalcification. The hypothesis is then accepted or rejected after an evaluation based on a set of criteria obtained from the models.

1.6 Objectives

The purpose of this research is to develop a complete procedure for detecting clusters of microcalcifications, beginning with a digitized mammogram and ending with the coordinates of the ROIs that contain the clusters. Also provided is a set of features that indicates the likelihood that the identified ROIs contain a cluster of microcalcifications. This purpose will be met by accomplishing the following objectives: 1) demonstrate the effectiveness of a wavelet based filtering scheme for segmenting microcalcifications, 2) present new techniques for isolating and new features for identifying individual objects in a ROI, and 3) introduce a new means of reducing the non-stationary portion of the background within a ROI and show how this improves the detectability of microcalcifications.

1.7 Organization

Chapter 2 contains a summary of past and current efforts to solve the breast cancer detection problem and an explanation of the theories used to develop the algorithms that have been implemented to detect breast cancer. Chapter 3 describes the algorithms and the process of using them. Chapter 4 gives the results, and Chapter 5 contains some concluding remarks and recommendations. Three appendices are also included which contain the specific code and some more detailed explanation of how the algorithms were developed and used.

II. Background

2.1 Introduction

This chapter reviews published research findings and theoretical developments related to the following three topic areas:

- Multiresolution/Wavelets
- Decision Boundary Feature Analysis
- Model-Based Vision

These three topics constitute the main tools used in this research. Wavelets have found broad usage in a number of engineering applications, and have been successfully applied to digital mammography in particular. Feature analysis based on decision boundaries and model based vision are also general pattern recognition techniques, but as far as can be determined, they have not been applied to digital mammography.

2.2 Multiresolution/Wavelets

Multiresolution analysis is a widely used technique in signal processing that incorporates the decomposition of the signal of interest into specified frequency bands (or, in wavelet terminology, into different scales). The decomposition of the signal into such separate bands is usually accomplished by a quadrature mirror filter bank (QMF). The filters in QMF banks are specifically designed to have certain properties that produce the desired effects of alias cancelation and perfect reconstruction[36].

Filters derived from the wavelet research of Mallat and Daubechies can be considered a special subset of QMF filters that satisfy certain mathematical properties [6] [24]. Wavelet filters have been shown to have several advantages over more traditional frequency analysis techniques, such as the Short-Time Fourier Transform[30]. These advantages are especially useful when analyzing non-stationary signals, such as digital mammograms. The decomposition of a two-dimensional image into its subband images via one dimensional filters is usually accomplished using a tree structured filter bank, as demonstrated by Jawerth[16]. This decomposition is traditionally used for image compression; however, it has also been shown to be useful for feature extraction as well[4].

Yoshida *et al.*[38] have shown that a tree structured filter bank based on the family of Least Asymmetric Daubechies (LAD)[6] wavelets can be successfully used for the automated detection of clustered microcalcifications. Their work demonstrated that reconstructing the mammogram from only the second and third level horizontal, vertical, and diagonal detail coefficients significantly improved the contrast between the microcalcifications and the background. Thresholding and texture analysis techniques were then used to identify specific 5.1-by-5.1 mm areas that were likely to contain at least three microcalcifications. Their preliminary results using a database of 39 mammograms with 41 microcalcification clusters yielded a detection rate of 85%, with a false positive rate of 5 clusters per image.

Laine *et al.* [20][21][19] have done similar research using separable one dimensional dyadic wavelets, a two dimensional ϕ -transform, non-separable hexagonal wavelets, and Deslauriers-Dubuc interpolation wavelets. A multiresolution decomposition was computed for each transform method, with a multiscale edge set of coefficients being created to highlight the desired image features. The edge sets were created along 45 degree axes for the dyadic, ϕ -transform, and Deslauriers-Dubuc coefficients, and along 60 degree axes for the hexagonal coefficients. The coefficients were processed using an adaptive threshold and gain which increased the signal to noise ratio between the microcalcifications and the surrounding tissue without increasing the signal to noise ratio between background artifacts and the surrounding tissue. The edge sets were then used to reconstruct an enhanced mammogram image. The work emphasized visually enhancing an image prior to examination by a human expert, rather than finding specific locations within an image that may be cancerous. Therefore, the researchers did not report results in terms of percentages of missed or detected microcalcification clusters. Instead, a contrast improvement index ($CII = \frac{C_{processed}}{C_{original}}$) for each method was calculated, using the definition of contrast introduced by Morrow[25]. Results varied between the methods depending upon the type of cancer that was being detected (small calcification cluster, spicular lesion, etc.). All four methods produced better contrast improvement than other traditional enhancement techniques, such as histogram equalization, with the Deslauriers-Dubuc interpolation based methods yielding contrast improvement indices of greater than ten for various microcalcification formations.

Strickland and Hahn[34] use wavelets as a bank of matched filters to detect microcalcifications. Modeling the background breast tissue with a non-stationary mean and a residual stationary Markov process, they derive the optimum pre-whitening matched filter for detecting a deterministic microcal-

cification. However, to compensate for their random nature the microcalcifications are modeled using a simple Gaussian form, and a bank of matched filters is derived based on the underlying assumptions of the statistical properties of the calcifications and the background. In their work, they show that the Gaussian object / Markov model for highly correlated noise matched filter corresponds closely to a biorthogonal wavelet using the spline variant as proposed by Cohen *et al.*[5]. A four octave filter bank is employed, with three intermediate voices between the octaves to give uniform coverage over the whole range of expected calcification sizes. The FROC curve generated by the method produced results comparable to other researchers, such as Karssemeijer[17].

Clarke *et al.* [27][28][29] have used algorithms that employ both two and three channel tree structured filter banks to detect and extract clusters of microcalcifications. In their research, they demonstrate the increased effectiveness of a three channel QMF bank over a two channel. The QMF filters are designed to meet the perfect reconstruction (PR) and linear phase criteria. Digital Mammograms are first pre-processed using a bank of tree-structured median filters which enhance the contrast of the microcalcifications and suppress background noise. These filters have a demonstrated superiority over standard median filters for retaining desired features without sacrificing noise reduction. The mammogram is then decomposed using the QMF bank and reconstructed using only a subset of the channels which were shown to contain most of the cancer data. A multistage artificial neural network using a backpropagation algorithm with Kalman filtering was used to detect microcalcification clusters. Using a set of 30 images, with 10 containing normal tissue and 20 containing 21 biopsy proven microcalcification clusters, they achieved a 90% detection with an average of .85 false positive clusters per image. Without preprocessing the results were a 68% true detection rate with 1.43 false positive clusters per image. When applied to a set containing 15 biopsy proven microcalcification clusters superimposed over simulated images they achieved a 100% detection rate with an average of .6 false positive clusters per image.

2.3 Feature Extraction

Feature selection is an important operation in pattern recognition. The idea is to find a set of vectors that represent an observation while reducing the dimensionality of the data. For pattern recognition, it is also desirable that the extracted features provide the best class discrimination. In other words, the goal is to find a set of basis vectors which form a subspace of the original data space;

when the data is projected into this new subspace, it still has the same class separability as it had in the original higher dimensional space.

Lee and Landgrebe[23] propose a method for linear feature extraction based on decision boundaries. Using a training set of pre-classified data, Bayes' decision rule is used to compute the location of the decision boundary based on an assumed Gaussian pdf of the data. The decision boundary is defined as the set of points in the N dimensional feature space where the probability of a data point being a member of class 1 or class 2 is equal. In other words,

$$\ln \frac{P(\omega_1)}{P(\omega_2)} = 0.$$

Several points of the decision boundary over a significant portion of it (referred to as the "effective decision boundary") are identified, and then the vector normal to the decision boundary at each of these points is computed. The outer product of each of these normal vectors are added together and averaged to form the effective decision boundary feature matrix Σ_{EDBFM} . Lee and Landgrebe show that the eigenvectors corresponding to the significant eigenvalues of Σ_{EDBFM} represent the "informative features", i.e., linear combinations of features that are optimum for preserving class separability. The space spanned by the orthogonal informative features forms a subspace into which the original data can then be projected while maintaining the same class discrimination that existed in the original space.

To solve the problem of non Gaussian data, Lee and Landgrebe[22] propose a similar method for feature extraction using non-parametric classification. Points along the decision boundary are identified using a non-parametric pdf estimate such as the K-nearest neighbors or the Parzen window methods. Once points along the decision boundary have been identified, the normal to the decision boundary at each point is calculated, and Σ_{EDBFM} is generated as before. This method is generally more accurate for data that is not well approximated by a Gaussian pdf, however, large data sets are often needed to accurately represent the effective decision boundary.

Feature saliency relates closely to feature extraction. The idea is to identify which features contribute most to class discrimination. Eisenbees[11] expanded on the basic feature extraction algorithm proposed by Lee and Landgrebe by examining the individual components of the dominant eigenvectors of Σ_{EDBFM} . He showed that there is some correlation between the individual features and their contribution to class discrimination based on how large their components were in the eigenvectors

which corresponded to the largest eigenvalues. Using a subset of the feature set based on this evaluation, Eisenbees was able to reduce his feature set from 256 to 30 and still preserve an overall classification accuracy of 94%.

Stewart[33] also incorporated the underlying idea of Lee and Landgrebe to identify salient features. Using a neural network to compute the decision boundary, Stewart derived a method to compute the eigenvalues of Σ_{EDBFM} from the gradient of the net's discriminant function. In this manner salient features can be identified from the interrelationships between the weights in the network. Stewart was able to derive a closed form solution for finding Σ_{EDBFM} , and thus identify the salient features.

2.4 Model-Based Vision

Model-Based Vision[2] is a pattern recognition scheme that uses previously derived models of items of interest to test a hypothesis of whether these items are present in an image. The scheme is usually divided into modules or stages: these are focus of attention (FOA), indexing, and matching. Focus of attention is intended to segment the image by identifying regions where there is a high likelihood of a target being present, and eliminating all other regions from further consideration. Indexing extracts features to make a decision as to what type of target may be present, and/or rank order the likelihood of specific targets being present. Matching compares the features extracted from the object in the image with features derived from a model that was created to simulate the target within the feature space as close as possible. The matching features are evaluated according to a set of criteria; if the features meet the criteria, then the object is considered to match that type of target. Model based vision has been traditionally used in military applications such as target identification, however its usefulness in cancer detection will be demonstrated.

2.5 Conclusion

Multiresolution analysis, decision boundary based feature selection, and model based vision target recognition are the underlying tools and methods that constitute the cancer detection procedure. Model based vision is the framework within which the detection procedure operates, multiresolution analysis using wavelets is the main focus of attention tool, and feature analysis will be used to make

a matching determination. The next chapter gives a detailed explanation as to how these tools are implemented.

III. Methods

3.1 Introduction

Having described previous research efforts, this chapter will now explain the actual methods used to distinguish cancerous from non-cancerous tissue regions.

3.2 Database

The nature of the data used for cancer research is of utmost importance in terms of being able to correctly evaluate the methods and results. Data for this thesis was obtained from the Wright-Patterson Medical Center. Dr. Jeffrey Hoffmeister, from Armstrong Lab, served as the medical expert for determining the cancerous condition of the cases that were evaluated.

The data consisted of 53 mammogram x-ray images that contained radiologist identified microcalcification clusters. The images were produced by a Loral film digitization system. The x-ray films were scanned and digitized to 12 bits (4096 gray levels) at a resolution of 100 microns per pixel using a Lumiscan 200 laser film digitizer interfaced to a Macintosh computer. The digitized data was processed on a Sun SparcStation computer.

3.3 Method Overview

The entire end-to-end procedure for detecting microcalcifications is shown in Figure 1. Subsequent sections of this thesis will explain each step in detail. Briefly, the procedure begins with the digitized mammogram which is processed by a non-linear operator to improve the dynamic range of those gray scale values which contain the most relevant information. The image is filtered by a wavelet transform, and then reconstructed using select bands. The image at this point is referred to as the *filtered image*. The filtered image is then thresholded and sub-divided into non-overlapping windows. Regions of interest (ROIs) are extracted from the filtered image by first allowing the windows to "settle" on areas where the local energy is maximum; the energy in each window is calculated with redundant windows removed. The coordinates of those windows which contain at least a certain minimum energy level, as determined by the statistics of the filtered image, become pointers used to extract ROIs from both the filtered and the original images. A high-pass operator subtracts out the low frequency non-stationary structure of the original ROI. Both sets of ROIs are then subjected to another thresholding routine that

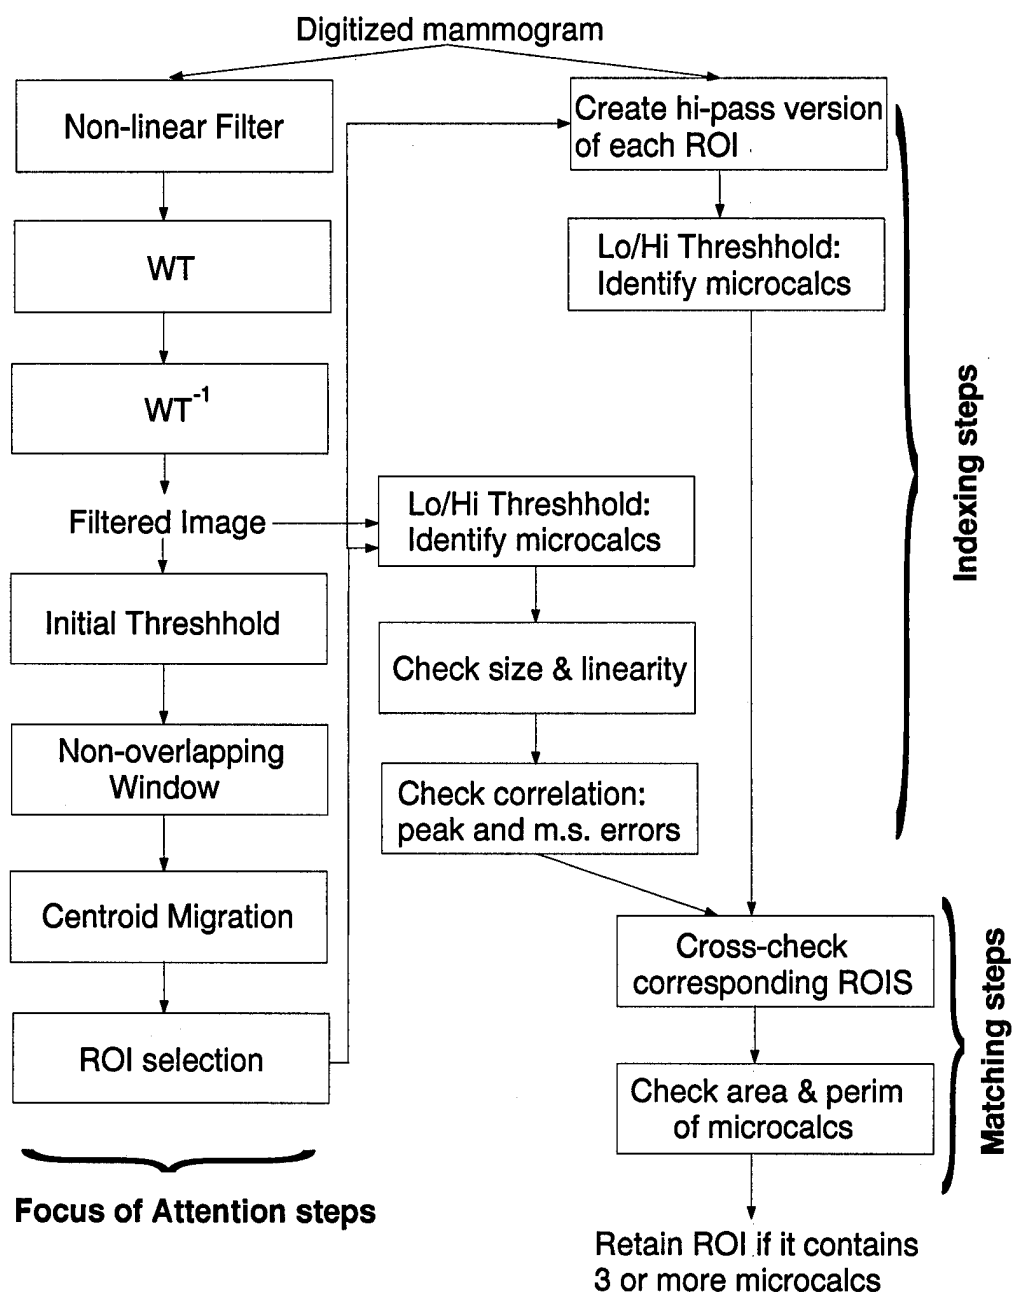


Figure 1. Entire procedure for detecting microcalcification clusters, starting with the digitized mammogram and ending with a set of ROIs


identifies objects within the ROI that may be individual microcalcifications. Features for each such object in the filtered image are extracted and evaluated; objects which survive feature examination are compared with their corresponding objects from the high-pass filtered original mammogram. If they match, and they meet a set of area and perimeter criteria, the object is considered a microcalcification. If an ROI contains at least three such microcalcifications, it is considered valid and passed along as a detection.

3.4 *Focus of Attention*


The focus of attention (FOA) portion of the algorithm consists of the steps starting with the original digitized mammogram and ending with the identification of ROIs that are likely to contain microcalcifications. This is analogous to a target identification process [2], which takes in raw image data and pertinent sensor information and outputs ROIs, along with confidence measures and characterizations of the image in and around the target. The steps constituting the FOA procedure are explained below. The goal of this procedure is to reduce the amount of data to be processed in the next step by 80 to 90 percent, without eliminating actual ROIs.

3.4.1 *Gray Level Modification.* The first step is to re-scale the gray levels of the mammogram such that the dynamic range of the information portion of the image is increased. This is done by applying a non-linear function to the pixel values. The original data was scanned in using a 12 bit linear gray scale (0-4095). It was determined from observation that most of the information pertaining to microcalcifications and parenchymal tissue was contained within the range of values from 2200 to 3600, which occupies only about 29% of the total available resolution. Figure 2 (a) shows a sample digitized mammogram and (b) it's accompanying histogram, with the hump of the histogram centered near 3000. A sigmoidal function was designed that expanded the values between these points so they have a higher dynamic range (i.e., the values from 2200 to 3600 are stretched to occupy a disproportionately larger portion of the scale from 0 to 4095) - see Figure 2 (b). In addition to increasing the dynamic range, this also has the desirable effect of providing a modest contrast enhancement between the microcalcifications and the background tissue.


For example, consider Figure 3 (a) and (b) which shows a 64 by 64 pixel ROI containing a small cluster of microcalcifications. Figure 3 (a) shows the ROI with its original gray scale values,



(a)



(b)



(c)

Figure 2. (a) Original digitized mammogram. (b) Histogram of the gray levels for the mammogram in (a) - the spike near 800 corresponds to the dark background around the breast, most of the relevant information is in the range from 2200 to 3600. (c) Comparison between the linear and non-linear relationship of pixel value to gray-scale value. For the non-linear relationship, the pixels in the steepest portion of the curve have the highest dynamic range and contrast.




Figure 3. 64-by-64 pixel ROI containing microcalcifications. (a) Original (unmodified) pixel values (b) rescaled ROI using the non-linearity showing increased contrast

and (b) shows the same ROI after the non-linear function has been applied to it. The contrast has been visibly improved; however, a more quantifiable measure of improvement can be computed. Using the definition of contrast introduced by Morrow[25], in which the contrast C is defined as

$$C = \frac{f - b}{f + b}, \quad (1)$$

where f is the mean gray level of the particular object in the image, and b is the mean gray level of the background pixels around the object. The values for f and b were calculated from the gray levels of the microcalcifications and the background, respectively. A mask was used to separate the microcalcifications from the rest of the image. Table 1 shows the improvement in dynamic range and contrast after applying the non-linearity. CII is the contrast improvement index introduced by Laine[21]:

$$CII = \frac{C_{processed}}{C_{original}}. \quad (2)$$

For this thesis the dynamic range improvement index (DRII) is similarly defined as:

ROI measurement	Contrast	CII	Dynamic Range	DRII
Original - mean	.0463	-	745.2	-
standard dev.	.0236	-	253.3	-
Scaled - mean	.206	4.25	1733.3	2.37
standard dev.	.125	.921	552.8	.425

Table 1. Table showing the contrast and dynamic range improvement produced by using a non-linear function to re-scale the gray levels. Values were obtained over a sample space of fourteen images, see Equations 2 and 3 for definitions of CII and DRII.

$$DRII = \frac{DR_{processed}}{DR_{original}}, \quad (3)$$

where the dynamic range is defined as the maximum gray value in the image minus the minimum gray value.

The table shows a mean contrast improvement of 4.25 and a mean dynamic range improvement of 2.37. Using global enhancement techniques, Laine achieved contrast improvement indices of 6.2, 6.4 and 5.8 for microcalcification clusters. Although the gains achieved by applying a non-linearity are more modest, they are much simpler to implement and are used only as a pre-processor. The overall objective of pre-processing is to enable the subsequent filtering to more easily separate the microcalcifications from the background.

3.4.2 Filtering. After modifying the gray level values, the image is then filtered by a tree structured wavelet decomposition/reconstruction scheme. A diagram of the basic method for 1 level of decomposition of a two dimensional image via one dimensional filters is shown in Figure 4; this is identical to the structure described by Jawerth *et al.*[16], and is widely used in filtering applications[38, 6, 30, 4].

Let $f(x, y)$ represent the input image, which is decomposed into the approximation image $f_{LL}(x, y)$ and the three detail images $f_{LH}(x, y)$, $f_{HL}(x, y)$ and $f_{HH}(x, y)$. $\tilde{H}(x)$ and $H(x)$ represent forward and reverse low pass filtering in the x direction; similarly, $\tilde{G}(y)$ and $G(y)$ represent forward and reverse high pass filtering in the y direction. $2 \downarrow x$ and $2 \downarrow y$ represent decimation by 2 in the x and y directions, and $2 \uparrow x$ and $2 \uparrow y$ represent upsampling (inserting a row or column of zeros) in the

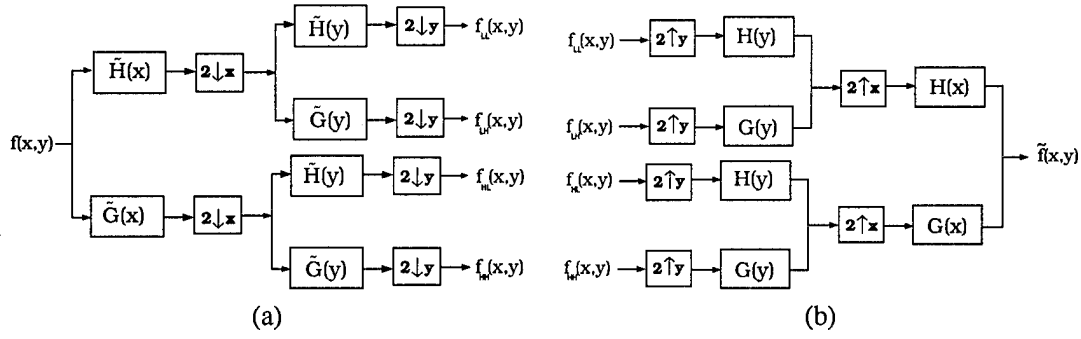


Figure 4. Block diagram of one level of (a) wavelet decomposition and (b) wavelet reconstruction [16].

x and y directions. If a perfect reconstruction scheme is used, then the output $\tilde{f}(x, y)$ will differ from $f(x, y)$ only by a delay d and possibly a scale factor c ; therefore, $\tilde{f}(x, y) = cf(x, y)z^{-d}$.

Although virtually any wavelet may be used in theory, the LAD family was chosen due to its regularity properties. It has been suggested[6] that regularity is more important for use in image processing, in terms of accurately representing rapidly changing structures, than other considerations, such as the number of vanishing moments. The LAD wavelets have also been shown to work well with mammograms in particular[38]. The 12 tap LAD wavelet (LAD12) was selected for use as the basis for filtering.

As indicated above, the image is decomposed into an approximation image and three detail images; each of these sub-images is $\frac{1}{4}$ the size of the original, and are usually represented in a Figure similar to 5 (a). The “1” subscript indicates that each sub-image is at a scale 1 below the original. The decomposition process can then be repeated on any of the sub-images, further decomposing it into an approximation and three detail images. Although in most cases, only the approximation image is further split into successively narrower bands, as shown in Figure 5 (b).

Figure 6 shows a mammogram that has been decomposed down to the level three approximation image and levels 1, 2, and 3 detail images. The reconstruction is carried out using only the coefficients from the level two and three approximation images. The other coefficients are all set to zero. In effect this is very similar to bandpass filtering the image with a passband of $\frac{\pi}{8}$ to $\frac{\pi}{2}$. This band of frequencies correspond to transitions or bumps that are roughly 2 to 8 pixels in size. The digitized mammograms have a spatial resolution of .1 mm per pixel, so 2 to 8 pixels is close to the expected size of microcalcifications. Giger *et al.*[38], Clarke[28], and Laine[21] have shown good results using




Figure 5. Wavelet decomposition to (a) first level (b) third level.

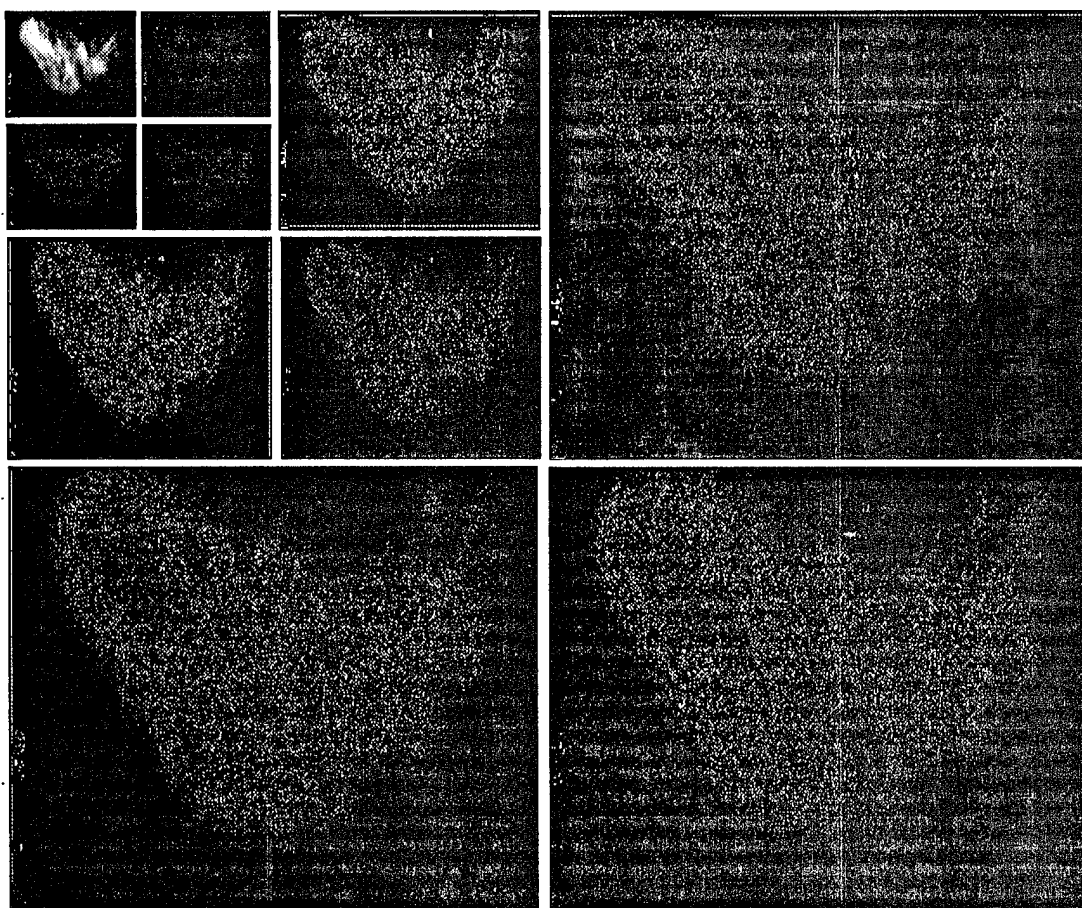


Figure 6. Mammogram that is wavelet decomposed to third level.

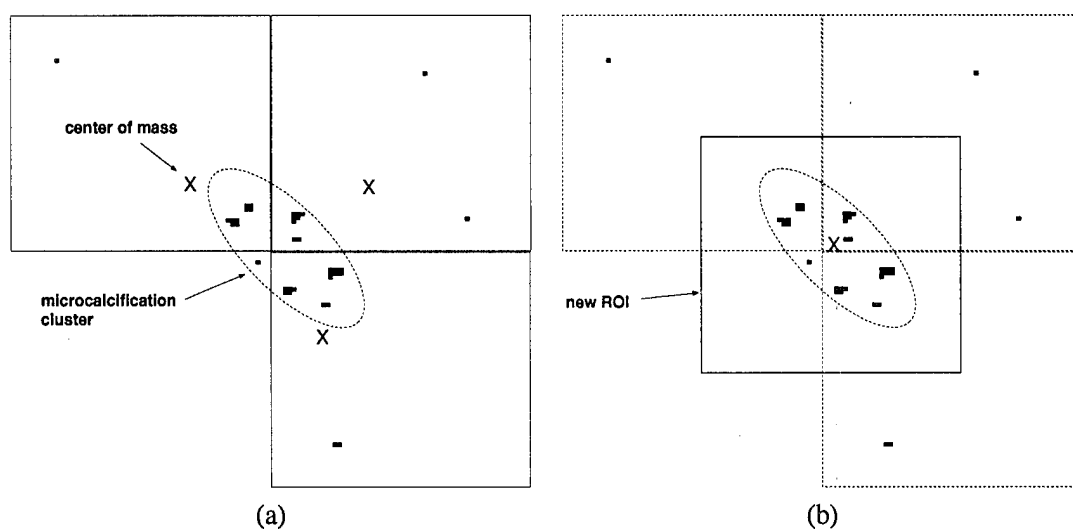


Figure 7. Demonstration of centroid migration (a) three separate windows each contain a portion of the microcalcification cluster; (b) final ROI after several successive centroid migrations containing the entire cluster.

the same or similar frequency bands. The reconstructed image is then clipped at a gray level of 255 and thresholded at zero (negative gray scale values are set to zero). This forms the filtered image from which ROIs are extracted.

3.4.3 Region of Interest Identification. Regions which potentially contain microcalcifications are found by identifying small boxes within the image (64 by 64 pixels) which have a relatively large number of bright pixels compared to most other same sized regions. The filtered image is thresholded at $T = 6\frac{1}{2}\sigma$, where σ is the standard deviation of the gray levels of the entire image. The value $6\frac{1}{2}$ was selected after a ROC analysis of the system performance on the training set, as discussed in Chapter 5. Then the 1020 by 2028 pixel image is subdivided into 465 non-overlapping 64 by 64 pixel windows (15 windows high by 31 windows wide, minus a small border around the edge). The sum of all the pixel values in each window is computed, which is used as a measure of energy (since there are no negative values); those with an energy of $4T$ or greater are retained for further consideration. This threshold value was chosen to reflect that microcalcifications are among the brightest objects in the reconstructed image, and a single microcalcification will be at least 4 pixels large. Therefore, any window containing at least one microcalcification will meet this criteria.

The center of mass for each retained window is then calculated, and the window is then re-centered around its centroid. This process is then repeated until the centroid moves less than 2 pixels, at which time it is considered to have “rested” in a region where it encloses the most energy. This process of “centroid migration” will tend to pull neighboring windows together when they each may have originally enclosed a portion of a cluster of microcalcifications. Once each of the original windows have finished migrating, they will most likely overlap each other (i.e., their centers of mass will be very close). Overlapping windows can be considered redundant if their centers of mass are within $d = 20$ pixels of one another, where $d = \sqrt{\Delta_{row}^2 + \Delta_{col}^2}$; the window with the lowest energy is eliminated. An example of this process is illustrated in Figure 7. If a cluster of microcalcifications happens to lie in part of three windows, and there isn’t a lot of other energy in the windows (usually a reasonable assumption), then the centers of mass (indicated by the X’s) of the three windows will all be near the cluster. As each of the windows migrates toward the cluster’s center of mass, the three windows will increasingly overlap until they are close enough to be considered redundant; then only one window will be retained.

After the original windows have migrated and the redundant windows have been removed, the energy of the new set of windows is calculated. Only those with energy equal to or greater than $12T$ are retained. This energy level roughly corresponds to the potential existence of at least three microcalcifications within the ROI. Those ROI’s which survive this process will be more closely scrutinized during the indexing procedure to determine if they do contain microcalcifications or background tissue that correlated well with the passband of the wavelet filter.

3.5 Indexing

The purpose of indexing is to formulate an initial hypothesis regarding each ROI which is then subsequently tested in the matching stage[2]. For microcalcification detection, the hypothesis is a simple yes or no decision on whether the ROI warrants further investigation. The goal of this step is to eliminate those ROIs which possessed enough energy to survive the Focus of Attention stage, but whose structure (as determined by an evaluation of several features) is more likely to correspond to background tissue than microcalcifications.




Figure 8. Cluster identification within an ROI: (a) Original filtered ROI, (b) ROI after thresholding with a high value, (c) ROI after thresholding with a lower value, (d) result after several iterations of morphological dilation of the the ROI in (b) and logical ANDing with the ROI in (c). The remaining four objects are then evaluated to determine their likelihood of being microcalcifications.

3.5.1 Index Features. Each ROI examined in this stage has a set of features generated which describe every object within it. These features and how they are generated are described in the following paragraphs.

The first step is to identify the individual objects within the ROI as illustrated by Figure 8. Two masks of the ROI are generated by thresholding the filtered ROI with a high value for one and a lower value for the other. Any objects within the ROI that consist of only one or two pixels are eliminated. The high threshold mask is then morphologically dilated (without letting any two objects grow together) and logical AND compared pixel by pixel with the lower threshold mask. This process

is repeated iteratively until the objects stop growing. The remaining objects are considered to constitute the potential microcalcifications (an object is defined as a group of pixels that are connected to one another by a continuous path of horizontal or vertical neighboring pixels). Each of the objects is then examined individually to determine its likelihood of being a microcalcification.

The first feature of an object is its size: an object must be between 4 and 50 pixels in size to be considered a potential microcalcification. The next feature is linearity, which describes the degree to which an object resembles a straight line. Strong linearity is associated with background structure (arteries, veins, etc.), whereas weak linearity indicates that the object has a blob shape, and is therefore more likely to be a microcalcification. The object's linearity is based on the Hough transform as described by Gonzalez and Wintz[14]. The linearity of an object at angle θ is

$$\gamma(\theta) = \frac{1}{l} \sum_{i=1}^l (i_x \cos(\theta) + i_y \sin(\theta))^2 \quad (4)$$

where i_x and i_y are the row and column coordinates of pixel number i in the object. The linearity can be generated from a matrix Ψ :

$$\Psi = \begin{bmatrix} P_{\theta_1,1} & P_{\theta_2,1} & \cdots & P_{\theta_n,1} \\ P_{\theta_1,2} & P_{\theta_2,2} & \cdots & P_{\theta_n,2} \\ \vdots & \vdots & \ddots & \vdots \\ P_{\theta_1,l} & P_{\theta_2,l} & \cdots & P_{\theta_n,l} \end{bmatrix} \quad (5)$$

where $P_{\theta_n,i} = i_x \cos(\theta_n) + i_y \sin(\theta_n)$. Let $\vec{\sigma}$ be a vector consisting of the standard deviations of each column of Ψ . Each entry in $\vec{\sigma}$ corresponds to the standard deviation of the linearity of the pixels at a specific angle from $-\pi$ to π . Therefore, the position of the smallest value in $\vec{\sigma}$ corresponds to the angle of maximum linearity of the object. The object's linearity L is therefore defined as

$$L = (\min(\vec{\sigma}))^2 / l \quad (6)$$

where l is the number of pixels in the object and $\vec{\sigma}$ is the vector of standard deviations. The lower the linearity factor, the more closely the object approximates a straight line (an object that is a straight line has a linearity of zero).

The other features represent the amount of correlation that exists between the filtered and the original images. Microcalcifications can typically be modeled as bright specks with some size distribution on a darker, noisy background[34]. Ideally, the objects identified in the ROIs from the filtered image (see Figure 8) exactly correspond to microcalcifications. In practice, however, this is not the case. Many objects do not correspond to microcalcifications at all, in which case they are considered to be some type of background tissue. If an object in a filtered ROI represents an actual microcalcification, then the pixels corresponding to that object in the original image will (in general) be the brightest pixels within the local area immediately surrounding the object (defining the local area to be the pixels within a 3 or 4 pixel wide border around the object). On the other hand, if the object does not represent a microcalcification, then it is likely to be part of an edge (which is why it was passed by the wavelet filter), and hence may not constitute the brightest pixels within its local area. To help distinguish microcalcifications from background which happens to meet the size and linearity criteria, a correlation method is used. Using the model of microcalcifications as bright specks on a low noise background, then the auto-correlation of the object in the filtered ROI (using a no-noise background) should approximate the cross-correlation of the object in the filtered ROI and its corresponding pixels in the original ROI. Two features are extracted and used to determine how close the two correlations are to each other. These are the center of the cross-correlation, which should be the highest value in the cross-correlation matrix, and the mean-square error between the two correlation matrices, which should be close to zero. This will separate microcalcifications from background tissue that may have had an edge that was passed by the filtering operation.

This procedure is demonstrated in Figure 9. An object representing a microcalcification is shown as a 3-dimensional mesh plot in (a). In (b) the mesh plot of the corresponding location in the original image is shown. The mean of the original ROI was subtracted out, and the image thresholded at zero to get the maximum contrast available within the ROI. The auto-correlation of the filtered object and the cross-correlation between the two images are computed and shown as mesh plots in (c) and (d). It can be seen that the two correlations are reasonably close to one another. The "closeness" is indicated by examining the height of the cross-correlation at the point where the auto-correlation is at its peak -

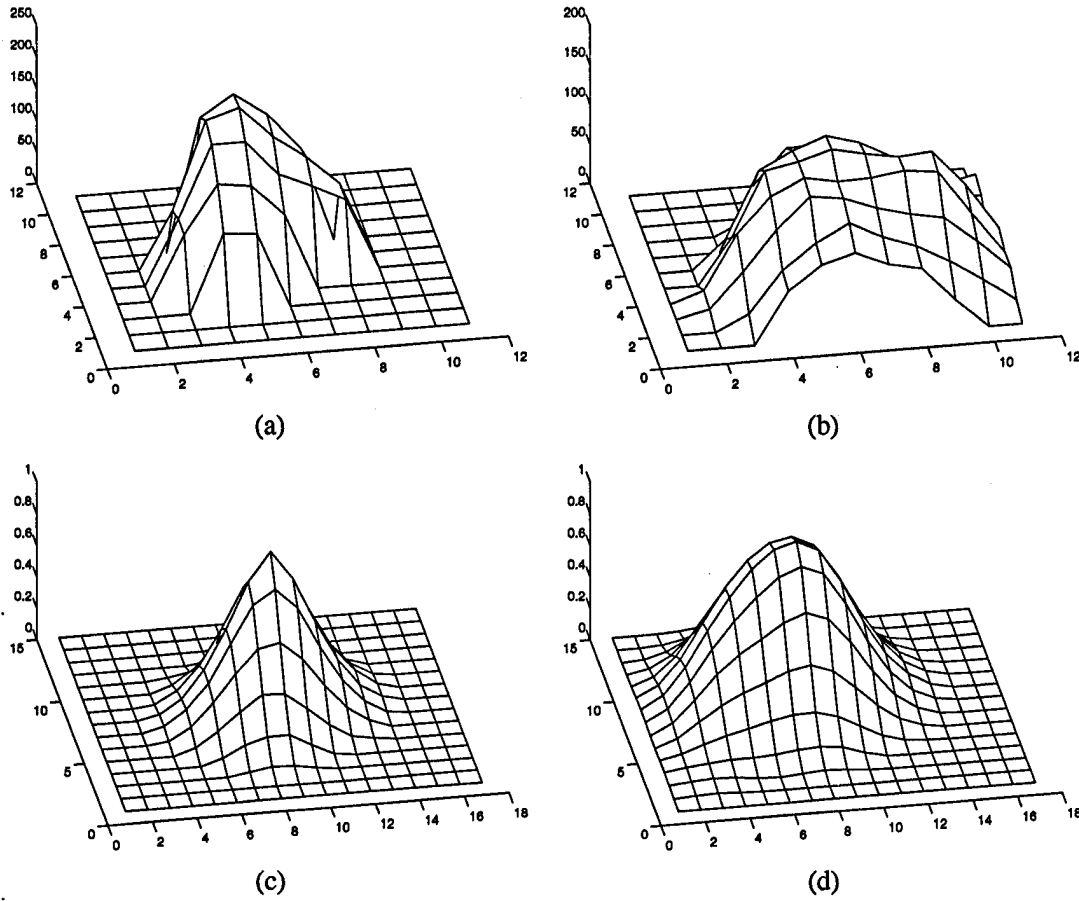


Figure 9. Mesh plots of (a) filtered microcalcification object (b) original microcalcification (c) auto-correlation of filtered microcalcification (d) cross-correlation of the filtered and original.

i.e. at the zero shift point. The other feature is the mean-square error. If \mathbf{R}_{AC} and \mathbf{R}_{CC} represent the N by M auto- and cross-correlation matrices, respectively, then the mean-square error ϵ is

$$\epsilon = \frac{1}{NM} \sum_{i=1}^N \sum_{j=1}^M (\mathbf{R}_{AC}(i, j) - \mathbf{R}_{CC}(i, j))^2. \quad (7)$$

Smaller epsilon is associated with the stronger the likelihood that the object is a microcalcification.

The converse is demonstrated in Figure 10. In this case the filtering operation passed the transition at the edge of some background tissue. The filtered object resembles a microcalcification

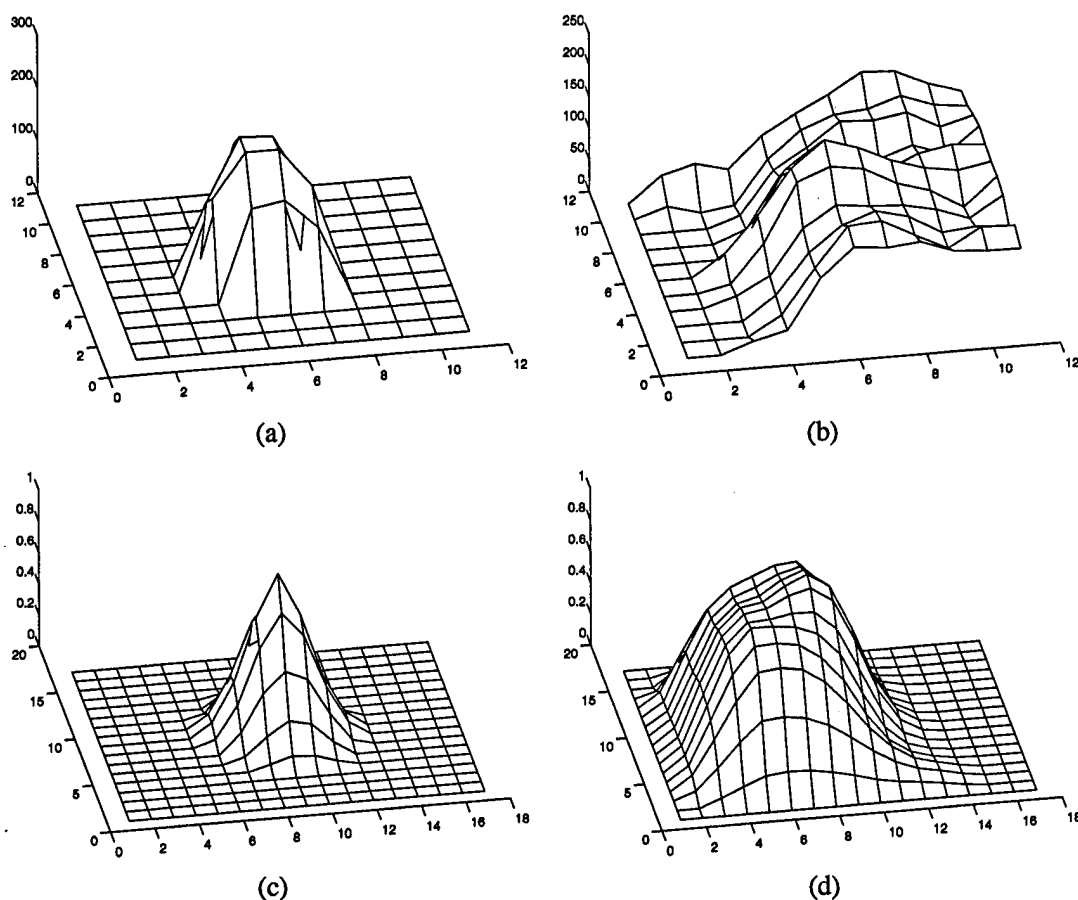


Figure 10. Mesh plots of (a) filtered "microcalcification" object (b) original image showing a tissue edge (c) auto-correlation of filtered microcalcification (d) cross-correlation of the filtered and original.

in size and shape, but the corresponding structure in the original image does not. Therefore, the mean-square error will be much higher for this case than the previous one.

3.5.2 Index Decision-making. The features described above are computed and used to make a determination as to whether the ROIs that survived the focus of attention stage are more likely to contain microcalcifications or background tissue. Each object in the ROI in question has its size and linearity computed. Microcalcifications should be at least four pixels and not more than 50 pixels in size. It was also determined by observation that a linearity of $L < 0.025$ is a good cutoff for separating microcalcifications (that are usually less linear and more "blob" shaped) and linear background tissue. Objects that meet these criteria are then evaluated using the correlation procedure to compute their

Feature	Criteria
size(S)	$50 > S > 4$
linearity(L)	$L > .025$
center correlation(C)	$C > .8$
mean-square error(ϵ)	$.11 > \epsilon$
# microcalcifications(N)	$N > 2$

Table 2. Feature decision thresholds and their order of consideration.

center correlation value and mean-square error. Ideally the center correlation should be one, but anything above .8 is retained. Also, in the ideal case the mean-square error is zero; however, it was determined that $\epsilon < .11$ is a reasonable cutoff. A ROI must have at least 3 objects that meet these criteria to be considered a possible microcalcification cluster. The decision steps are summarized in table 2.

3.6 Matching

3.6.1 Final Object Identification. In a typical target recognition application, the matching process tests a hypothesis about an object's identity by comparing features extracted from the object against some criteria. The criteria is usually based on a model which was developed to resemble the object in terms of the extracted features. The final step, then, in deciding whether a microcalcification cluster is detected is to extract features from each calcification and compare them with a model-based standard.

As previously discussed, microcalcifications can be modeled as bright specks on a non-stationary noisy background. Because microcalcifications exist over a very wide range of gray scale values, and the background tissue (referred to as "noise" hereafter) in which they exist also has a wide distribution of gray scales and amount of correlation, it is very difficult to identify a universal standard which can be used to separate microcalcifications from noise. Instead, it is more profitable to consider each object within a highly localized area consisting of the object itself and a few pixels immediately surrounding it; this allows for comparisons to be made on a relative basis. For this reason, it is helpful to have a second view of the object as a double-check of the attributes to be examined.




Figure 11. Comparison of object separation using the original ROI: (a) & (b) without removing the non-stationary background; (c) & (d) after removing the non-stationary background. Once the non-stationary portion of the background is removed, the dimmer microcalcifications are retained, and the background object on the bottom of the original image is discarded.

Such a second view is generated by creating a high-pass version of a ROI being considered, which removes the non-stationary portion of the background, leaving any microcalcifications on a zero-mean noisy background.

The non-stationary background is approximated by filtering the ROI with an 8-by-8 pixel averaging window. This effectively removes the signal portion of the ROI, and leaves a smooth approximation of the background structure. This structure is then subtracted from the original image, producing a high frequency only ROI consisting of microcalcifications and noise on a “flat” background. The hi-pass image can then be subjected to a thresholding scheme that is similar to the one used with the wavelet filtered image. The effect of this is demonstrated in Figure 11. Figure 11 (a) and (b) show the original

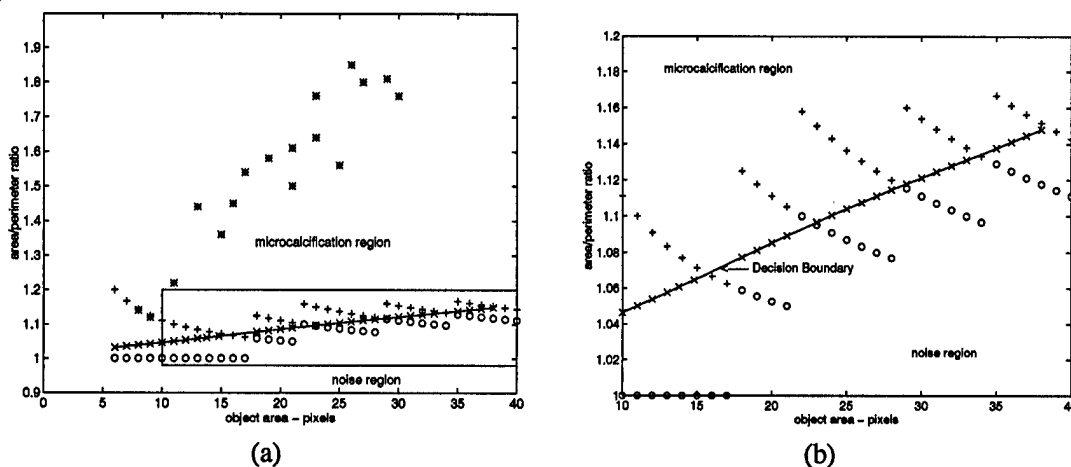


Figure 13. (a) Decision boundary and data points: * = phantom data points, + = actual microcalcifications, o's = noise data points. (b) Blow-up of the critical portion of the decision boundary and data points actually used to determine the decision boundary: + and o indicate data points, x's are decision boundary points as determined by Lee and Landgrebe's non-parametric procedure

3.6.2 Hypothesis Testing. Once the final ROI has been generated, objects within the ROI are subjected to the last test as to whether they represent microcalcifications or noise. The area and perimeter of each object is determined, and then the ratio of area/perimeter is calculated. If the ratio falls within a certain range based on the area, then the object is considered a microcalcification; otherwise it is considered noise. Figure 13 (a) illustrates most of the data region, including several sample points for both microcalcifications and noise, and an approximation to the decision boundary. An X-ray phantom image containing several modeled microcalcifications of assorted sizes was used as a guideline. The area/perimeter ratio was calculated and plotted for each modeled microcalcification. Since the phantom modeled microcalcifications are somewhat idealized in terms of their size and shape, they represent a high confidence level indicator. The (*) points represent modeled microcalcifications. The decision boundary was determined by selecting several examples of microcalcifications and noise, and generating an effective decision boundary feature matrix[23]. Figure 13 (b) shows a magnified view of the points used to determine the boundary. All + 's and o 's are actual data points. These points were used as input for the decision boundary calculation[22], and the x 's indicate the decision boundary points.

Each object in the final ROI is tested to determine which category it falls into - microcalcification or noise. Any ROI with at least three microcalcifications is passed along as being detected.

3.7 Conclusion

The detection procedure involves a number of different steps, each of which is designed to gradually eliminate more and more false ROIs. Actual results generated by this method will be presented in the following chapter.

IV. Results

4.1 Introduction

This chapter presents the results of applying each of the steps described in Chapter 3. In all, the procedure was applied to 53 mammograms containing a total of 53 ROIs with radiologist confirmed microcalcifications. The data was divided into four sets: a training set consisting of 14 mammograms with 17 ROIs, a test set consisting of 17 mammograms with 20 ROIs, an evaluation set consisting of 12 mammograms with 16 ROIs, and a set declared normal (i.e., cancer free) by a radiologist.

The training set was used to develop the procedures and methods discussed previously. Therefore, the results obtained from this set will receive more emphasis in the following discussion as its progress is traced through the steps in the detection method. Once this was done, the test and evaluation sets were processed and evaluated without changing the procedure.

4.2 Training Set

The training set consisted of 14 mammograms from 9 different patients. Of the 17 ROIs, 10 were confirmed as malignant after biopsy, 5 were confirmed benign after biopsy, and 2 were not biopsied.

4.2.1 Focus of Attention. The steps comprising the FOA portion of the method were applied to the images. After the images had been filtered by the wavelet filter bank and thresholded, the ROIs were ranked according to the amount of energy they contained. Table 3 shows the total number of ROIs that survived the FOA steps, and an energy rank of how the true ROI compared to all the others (1 = ROI having the highest energy in the image, etc.). Those images with more than one number in the rank column had more than one ROI, and were ranked accordingly.

The mean and standard deviation of both the number of ROIs passed and the rank order of the true ROIs is also listed in the table. The numbers in parentheses are calculated by ignoring image 008, which represented an extreme case in terms of the rank order of the ROI. The microcalcifications in image 008 are small and have low contrast, causing very little energy to be present in the segmentation steps. It could be argued that the number of false positives surviving the FOA stage could be substantially reduced if only the top 15 or 20 ROIs were retained, at the expense of losing only one true ROI. However, it was decided at this early phase of the procedure to allow the large number of ROIs to

Image	Num. Regions	rank
005	78	1
006	89	2,3,5
007	78	1,6
008	61	60
009	82	6
020	69	6
022	54	2
024	56	1
033	66	2
038	95	3
040	53	13
045	55	1
047	41	3
055	45	10
mean	66.4	7.35 (4.1)
σ	16.1	13.9 (3.4)

Table 3. Table showing the results of the focus of attention procedure. Num. regions indicates the number of ROIs that survived the thresholding procedure, and rank indicates the rank order by energy of the true ROI (1 = highest energy, etc.).

pass in exchange for passing all the true microcalcification clusters. Even though the number of ROIs passed is large, an average of 66 non-overlapping ROIs that are 64 by 64 pixels in size represent $\frac{66 \times 64^2}{1024 \times 2048} = .1289$ about 13% of the image, which means a reduction in the amount of data to be considered by about 87%, with all true ROIs being retained.

4.2.2 Indexing. Once the ROIs were identified, the next step was to make a decision whether they were more likely to represent microcalcifications or noise. The indexing steps evaluated the size, linearity, and degree of correlation to make the decision. Table 4 shows the amount of data that survived the indexing steps.

A comparison of Tables 3 and 4 shows that indexing removed an average of approximately 50 ROIs per image, while still retaining all of the microcalcification clusters. Again, if only the top fraction were retained, such as the top 4 or 5, a very good detection percentage could be achieved with a reasonable number of false positives. However, it was again decided to retain all of the ROIs that passed this step, and to use a matching scheme to make a final decision about each ROI.

Image	Num. Regions	rank
005	17	2
006	17	1,2,3
007	18	1,2
008	18	18
009	4	1
020	18	1
022	15	1
024	25	1
033	23	2
038	17	4
040	16	5
045	13	1
047	15	1
055	6	1
mean	15.8	2.76 (1.8)
σ	5.5	4.1 (1.2)

Table 4. Table showing the results of the indexing procedure. Mean and standard deviation numbers in parentheses are calculated by ignoring image 008.

4.2.3 Matching. The matching process compared the area and area/perimeter values for each potential microcalcification within an ROI against a standard that was developed from the phantom images. At its optimal setting, the method achieved a 100% (17 of 17) detection with an average of 1.64 false positive ROIs per image.

Since the performance of the method is dependent upon the settings of a number of threshold and comparison values, it is informative to examine how the performance changed as the parameters changed. Four receiver operating characteristics (ROC) curves are shown in figure 14 to demonstrate the effects of altering the parameters. The altered parameters are: (a) the multiple of the standard deviation used to determine the threshold value in the windowing of the wavelet filtered image, (b) the parameter controlling the low/hi threshold pair applied to ROIs that identified the separate objects within the ROI, (c) The correlation peak minimum and mean square error maximum values used to screen objects, and (d) the threshold parameter used to separate objects in the hi-pass ROIs prior to the cross check between the hi-pass and wavelet filtered ROIs.

4.2.4 Discussion. Examination of the ROC curves gives an indication as to how the performance of the method changes as the parameters are changed. The segmentation threshold

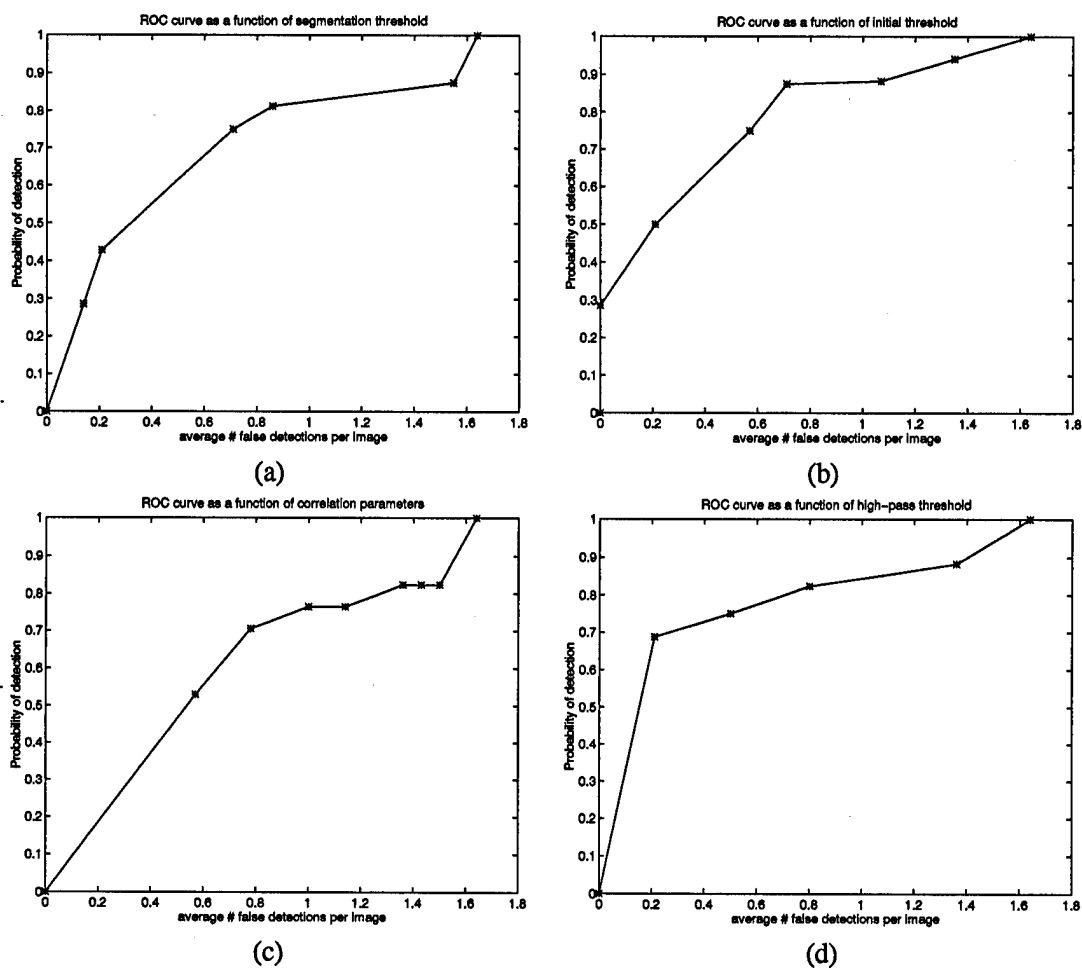


Figure 14. ROC curves showing performance as detection parameters are changed: (a) multiple of the standard deviation used for segmentation, (b) initial low-hi threshold pair, (c) correlation peak and mean square error criteria, (d) hi-pass ROI threshold

parameter (14 (a)) is the first hurdle that must be overcome in order for a ROI to be detected. The effect of raising this threshold is to eliminate lower energy ROIs, which causes ROIs containing few and or small microcalcification to be immediately discarded; false ROIs which contain low energy noise would probably be eliminated in one of the subsequent steps. This explains the steep initial drop in the curve as the ROIs corresponding to images 8 and 40 (see Table 3) are discarded, then the decrease becomes more gradual.

The low/hi threshold parameter (Figure 14 (b)) appears to offer the best tradeoff between decreasing probability of detection and the number of false detections. An explanation for this is that the ROIs are not evaluated in terms of their energy after this step, instead this is a pre-processing operation used prior to determining whether the individual objects within a ROI are more likely to be microcalcifications or noise. This means that the ROIs are retained or discarded based on the *distribution* of the energy within them, not necessarily on the *magnitude* of the energy. Therefore, proportionately more false ROIs are eliminated as the threshold increases as compared to the other parameters.

Changing the parameters of correlation peak and mean square error (14 (c)) appears to produce the most extreme changes at the ends of the curve. These are the only parameters that are not threshold values; instead they are feature criteria. In addition, both features are highly non-linear - particularly the peak correlation value. At its most extreme level, the peak correlation is a binary yes or no decision as to whether the correlation center is the maximum value in the cross correlation between the objects in the wavelet filtered image and the original image. Because of the non-linearity of these features, it is difficult to predict the effect that changing them will have on the method as a whole.

The hi-pass threshold parameter (14 (d)) governs the size and shape of the objects in the hi-pass ROI that are cross-checked with the objects in the wavelet filtered ROI. As the threshold increases, these objects get progressively smaller. This has a double-edged helpful effect on the detection process. Objects which are essentially the "wrong" shape to be microcalcifications (wrong being described as wispy, or not compact) are eliminated quicker than compact objects that more closely resemble microcalcifications, which are retained until they are too small to be detected at all. This also explains the steep drop at the end of the curve, which represents the threshold point beyond which very few microcalcifications survive the area/perimeter test.

4.3 Test Set

The procedure was used to evaluate the test data set using the same settings that were applied to the training set without modification. Fifteen of the twenty ROIs were detected (75% detection rate) with an average of 1.82 false positive detections per image. Applying the procedure without modifying any of the parameters produces a less "biased" estimate of its performance; the training set was used to develop the algorithm, therefore there is a tendency to customize portions of the method to match it. This becomes apparent when the results of the two methods are compared: 100% detection with 1.64 false ROIs per image for the training set versus 75% detection with 1.82 false ROIs per image for the test set. A more realistic estimate of how the method would work in a real world setting is obtained by using pristine data with the unmodified algorithm.

It is informative to examine the ROIs that were not detected and to discover the reason(s) why they were not. Two of the five microcalcification clusters were too small to meet the minimum energy requirement in the segmentation step. In the other three cases, all three clusters contained two calcifications that met all of the requirements, but one (or more) calcifications failed to meet some requirement by a narrow margin. Most missed microcalcifications could be attributed to one of two factors: either the correlation features were outside the acceptable limits, or the area/perimeter ratio was too low. The first reason for failure appeared to occur when the signal to noise level of the microcalcification was very low, which caused a strong cross-correlation between the noise and the microcalcification. The second reason appeared to be the result of having a threshold too high, causing too little of the microcalcification to be left from which the area/perimeter ratio was calculated. In either case, some additional pre-processing may help reduce the noise level such that a third calcification could be detected. Since an ROI must have three valid microcalcifications to be considered a detection, these ROIs were discarded.

4.4 Evaluation Set

The procedure was also applied to a third set of data, again without modifying any of the parameters. Twelve of sixteen ROIs were successfully detected, with an average of 4.25 false positive detections per image. The higher number of false positives was due to one case, which consisted of a pair of older images that were produced on different film. Excluding that particular case resulted in 8 of 12 detections with 2.8 false positive ROIs per image. For the missed ROIs, once more, most of

# Detections	# images which occurred
0	5
1	2
2	1
3	1
4	0
5+	1
total	10

Table 5. Number of (false) detections that were found by the algorithm in radiologist declared normal images.

them were attributed to cases where 2 microcalcifications were passed but a third narrowly failed some criteria.

4.5 Normal Tissue Set

As a final check, the procedure was applied to a set of 10 images that were declared by a radiologist to be free of abnormalities (i.e., no microcalcifications or masses were present). This experiment was intended to indicate how the algorithm would perform on a "normal" case by counting the number of detections. Table 5 shows the results of this test. One image produced 10 false positive detections giving an overall average of 1.7 detections per image. The reason for the single anomaly is unclear, however, it is suspected to be related to the variations in the density of the breast tissue. Visual examination of these images indicated that there may be some correlation between how homogeneous the breast x-ray appears and how likely a large number of false positive detections would occur.

4.6 Summary

For the training set the procedure achieved a 100% (17 of 17) detection rate, with an average of 1.64 false positive ROIs per image. When the three data sets are combined for a composite performance measure, 44 of 53 (83%) ROIs were detected with an average of 2.44 false detections per image. Combining only the test and evaluation sets resulted in 27 of 36 (75%) ROIs detected with 2.83 false positive detections per image. The composite results are summarized in the ROC plot in Figure 15. Combining the test, evaluation, and normal tissue results should give a reasonable estimate of the system performance for most mammograms.




Figure 15. Composite probability of detection and false positive ROIs per image results of training, test, and evaluation data

V. Conclusion

5.1 Introduction

This chapter summarizes the methods and results of this research, while highlighting the developments that are unique to AFTT and to cancer detection as a whole. The main purpose of this research was to produce a front to back procedure that begins with a digitized mammogram and outputs the coordinates of regions within the image that are likely to contain a cluster of microcalcifications. The procedure also outputs with each region a set of features to indicate how strong the likelihood is that the region contains a cluster of microcalcifications. The objective is to help a radiologist identify abnormal tissue, not to make a diagnosis; therefore, no attempt is made to classify abnormalities as malignant or benign.

5.2 Summary of Methodology

The steps involved in detecting ROIs are the result of combining some unique tools developed by this research with the application of concepts developed by other researchers that were outlined in Chapter 2. The first step is to apply a non-linear function to the image to improve the contrast and dynamic range of the most relevant information. Similar techniques have been used successfully by other researchers[20, 21, 12, 18]; for this research the AFTT non-linear transformation of the brightness values produced a mean contrast improvement index (CII) of 4.25 compared to 6.4 for Laine's more computationally intensive method[21, 19]. A new measure that is similar to the CII was also introduced in this work: a dynamic range improvement index (DRII), which measures how the relevant portion of the information can be spread over the available resolution in order to enhance detectability at the expense of shrinking the dynamic range of the less important information.

This research implements the wavelet decomposition/reconstruction method for image filtering demonstrated by Yoshida *et al.*[38]. The filtering increases the signal to noise ratio of microcalcifications, allowing for easier detection using thresholding and feature extraction. The intent was to test the applicability of Yoshida's filtering method to AFTT's mammogram database. Although the post-filtering portion of the method used by Yoshida to detect microcalcification clusters differed from what was used for this research, it is worthwhile to compare the results: Yoshida's complete procedure produced an 85% detection rate with an average of 5 false positives per image using a database of 39

mammograms, compared to this research which produced a training result of 100% detection with an average of 1.64 false positives per image and a test and evaluation result of 75% detection with 2.54 false ROIs per image using a total of 53 images. It is unclear whether or how Yoshida divided the 39 images into training or test sets.

After filtering and thresholding the image, the next step is to identify the center coordinates of ROIs that contain microcalcification clusters. A process called centroid migration is introduced which finds the center coordinates and eliminates redundant ROIs. This process, unique to this research, reduced the number of ROIs that were passed by the focus of attention portion of the algorithm by approximately 30%.

The background noise removal procedure was in part inspired by the work of Hahn and Strickland[34], who proposed modeling breast tissue as a stationary random process with some correlation on a non-stationary background. This research presents a novel approach to removing the non-stationary portion of the background by subtracting an estimate of the background which was created through mean filtering. Potential microcalcifications were then extracted using low/hi thresholding and morphology. The combination of non-stationary background noise removal and dual thresholding to extract microcalcifications is unique to this research, and removed on average 5 false positive ROIs per image from the training set.

Microcalcifications were confirmed as such by comparing features generated from each potential microcalcification with a model based criteria, using phantom microcalcifications as the modeling source. Although Clarke *et al.*[29, 28] and Giger *et al.*[12] have used phantom images to simulate microcalcifications, the extraction of features from phantoms for model based matching is unique to this research. New features of peak and mean-square error of the cross correlation between two estimates of microcalcifications are also proposed in this work; when combined with other standard features such as linearity, area, and perimeter it is shown that classification results comparable to those achieved by Dhawan[8], who using a neural network to evaluate features based on second order statistics to identify microcalcification clusters achieved a detection rate of 77% with a false positive rate of 45% on a database of 100 images.

5.3 Summary of Results

Set	# Images	True Detections	%	False detections	average/image
Train	14	17 of 17	100	23	1.64
Test	17	15 of 20	75	31	1.82
Eval	12	12 of 16	75	51	4.25
Clear	10	-	-	17	1.7
Total	53	44 of 53	83	122	2.3

Table 6. Detection and false positive results for all data sets.

The results for each of the data sets is shown in Table 6. Listed for each data set is the number of images in the set, the number of microcalcification clusters detected versus the number of clusters present, the percentage of clusters detected, the total number of false positive ROIs detected for the set, and the average number of false ROIs per image. The total for all of the data sets combined was 44 out of 53 ROIs detected (83%), with 122 false detections in 53 images, which yields an average of 2.3 false positives per image. It is important to notice the change in results for the different data sets. The training set was used to develop the algorithm and determine the various threshold parameters, indicating that it is possible to customize the algorithm to detect everything with a fairly low false ROI detection rate. The test, evaluation, and clear data sets were processed independently and with the exact same algorithm parameters as were used on the training set. Although the performance fell off, the results of the three independent sets were reasonably consistent with one another. From this it can be concluded that the same results could be expected from virtually any set of images.

The missed detections mainly occurred in cases of extremely small and/or poor contrast microcalcifications. In particular, low contrast calcifications most often failed to survive the initial thresholding of the wavelet filtered image. Small or high noise calcifications that were missed usually failed to meet one of the correlation feature criteria or the area/perimeter feature requirement. The missed detections are useful in that they define the boundaries of the system's performance; it is suggested that follow-on work be done to expand these boundaries by investigating other adaptive thresholding techniques in order to increase the detectability.

Although it was not the purpose of this research to diagnose detected ROIs as malignant or benign, it is informative to re-examine the detection results in terms of the number of malignant and benign cases detected. The 43 biopsied images comprise 27 different cases (16 double view and 11 single view), with 14 pathology confirmed malignant and 13 pathology confirmed benign. The system detected 13 of the 14 (93%) malignant cases with an average of 2.85 false detections per image. Of

the 14 malignant cases, 6 were used in the training set and the other 8 were distributed in the test and evaluation sets. For the 13 benign cases, 10 were detected (77%) with an average of 2.1 false ROIs per image. A case was considered detected if an ROI corresponding to a biopsied cluster in one of the views was output by the system; if the microcalcification cluster was not output for either view, it was counted as a missed detection.

It is interesting to note that the system was more successful in detecting malignant clusters than benign ones, and images with malignant clusters had a higher incidence of false positives. This indicates that the features used to separate microcalcifications from background tissue may potentially be used to separate malignant from benign clusters, as the system is more sensitive to malignant cases. It is also of practical interest: a missed benign cluster may not have serious repercussions for the survival of a patient but a missed malignancy could prove fatal. Two possible explanations of why a higher percentage of malignant cases are detected are: first, very small clusters that are susceptible to elimination during the initial filtering and thresholding steps are more likely to be benign and second, benign microcalcifications may be less separable from background tissue using area/perimeter as a feature than malignant ones.

5.4 Conclusion

Some interesting conclusions can be drawn from this work:

- Filtering the mammogram with a tree structured wavelet filter using a Daubechies wavelet increases the signal to noise level of the microcalcifications sufficient to allow a thresholding scheme to identify over 90% of the radiologist identified microcalcification clusters.
- Low and a high thresholding combined with morphological filtering proved successful in isolating individual microcalcifications 95% of the time. Features extracted from each calcification were then used to evaluate its authenticity, with the results as given above. Selection of the threshold values as a function of the statistics of the image was effective, but occasionally produced excessive false positive detections.
- It is possible to remove enough of the non-stationary background within an ROI through a simple mean filtering and subtraction technique to increase both the sensitivity and specificity of the process. For the set of training data this method removed an average of 5 false positive ROIs per image.

With these conclusions, this thesis meets the objectives outlined in Chapter 1. The procedure developed in this research may be applied to most mammograms and be expected to yield results similar to those found using the data in this research.

Appendix A. Wavelet Code

A.1 Introduction

This appendix describes the MATLAB m-files that were used to perform the wavelet decomposition and reconstruction of the mammogram images evaluated in this thesis. The wavelet code consists of two m-files: `new_wave.m` and `new_recon.m`. `New_wave.m` decomposes the input image into an approximation image and three detail images as described in chapter 3. `Recon_wave_conv.m` reconstructs the image from its four sub-images. Each of these is described in detail below.

A.2 Requirements

The image size is not constrained to be square, or even a power of 2. Because the filtering is done using MATLAB's `conv2` command, the filtered image size is equal to the original image size plus the filter length minus 1. That is, if the original image is $m \times n$, and the filters are of length l , the filtered image is of size $m + l - 1 \times n + l - 1$. The filtered image is then decimated by two so that each sub-image is $\frac{m+l-1}{2} \times \frac{n+l-1}{2}$. Generally, m , n , and l are even, so $p = m + l - 1$ and $q = n + l - 1$ are odd. However, the decimation routine automatically rounds $\frac{p}{2}$ and $\frac{q}{2}$ up to the nearest integers, a and b , so $a = \text{ceil}(\frac{m+l-1}{2})$ and $b = \text{ceil}(\frac{n+l-1}{2})$. The original image size $m \times n$ must therefore be chosen such that a and b are even for *all* of the decomposition levels that are to be computed. This is a function of the filter length, so a correct set of sizes must be computed for each filter length. For example, if the original image size is 2312×1160 , and the filter has 8 taps, the first level decomposition will be $2312 + 8 - 1 \div 2 = 1159.5 \rightarrow 1160$ by $1160 + 8 - 1 \div 2 = 583.5 \rightarrow 584$, or 1160×584 ; the second level will be 584×296 , etc. The reason for this non-linear size progression is the nature of the `conv2` command. Its disadvantages of awkward size requirements and non power-of-two progression are outweighed by its advantages of speed and minimizing the edge effects of filtering.

A.3 `new_wave.m`

`new_wave.m` is the routine that computes the wavelet coefficients for an image. It is a function that accepts an image and the wavelet filter coefficients for the low and highpass filters as inputs, and returns the approximation and detail images. The program is designed to be flexible so as to handle images of any size using any separable wavelet filter with the following restrictions: 1) The filter

coefficients must be the same size (i.e., it does not handle biorthogonal wavelets of different lengths),

2) The image must be a specific size if it is to be reconstructed (see above).

A wavelet decomposition is created by calling the function `wave_conv` for an image, which returns four sub-images. The next level of decomposition can then be created by calling the function again, only this time use the approximation image returned from the first call as the input image, and so on until enough decomposition has been achieved.

```
function [s11,slh,shl,shh] = new_wave(IMG,LO,HI);

%*****
%
% function [s11,slh,shl,shh] = new_wave(IMG,LO,HI);
%
% Uses the low and hi-pass filters LO and HI to
% create the approximation image s11 and the three
% detail images slh, shl, and shh from the original
% image IMG
%
%*****

% low-pass filter the columns
L1 = conv2(IMG,LO);

% downsample the columns
[nr,nc] = size(L1);
L1=L1(:,1:2:nc);

% low-pass filter the rows
s11 = conv2(L1,LO');

% downsample the rows
[nr,nc] = size(s11);
s11 = s11(1:2:nr,:);

% hi-pass filter the rows
slh = conv2(L1,HI');

% downsample the rows
[nr,nc] = size(slh);
slh = slh(1:2:nr,:);

% hi-pass filter the columns of the original
H1 = conv2(IMG,HI);

% downsample the columns
[nr,nc] = size(H1);
```

```

Hl = Hl(:,1:2:nc);

% low-pass filter the rows
shl = conv2(Hl,LO');

% downsample the rows
[nr,nc] = size(shl);
shl = shl(1:2:nr,:);

% hi-pass filter the rows
shh = conv2(Hl,HI');

% downsample the rows
[nr,nc] = size(shh);
shh = shh(1:2:nr,:);

```

A.4 new_recon.m

The program new_recon.m reconstructs an image from four sub-images and a set of filter coefficients that are provided as input.

```

function BACK = new_recon(sll,slh,shl,shh,LO,HI);

%*****
%
% function BACK = new_recon(sll,slh,shl,shh,LO,HI);
%
% Creates reconstructed image BACK from the approximation
% and detail images and lo and hi-pass filters LO and HI.
%
%*****

% upsample along y - i.e. double the number of rows:

[nrow,ncol] = size(sll); % all four must be the same size

sll_up(1:2:((2*nrow)-1),:) = sll;
slh_up(1:2:((2*nrow)-1),:) = slh;
shl_up(1:2:((2*nrow)-1),:) = shl;
shh_up(1:2:((2*nrow)-1),:) = shh;

% convolve the pieces with their appropriate low or high pass filter

i1l = conv2(sll_up,LO');
i1h = conv2(slh_up,HI');
ihl = conv2(shl_up,LO');
ihh = conv2(shh_up,HI');

```

```

% combine the intermediate images

[row,col] = size(ill);
inter_lo = ill + ilh;
inter_hi = ihl + ihh;

% upsample along x - i.e. double the number of columns:

intl(:,1:2:(2*col)-1) = inter_lo;
inth(:,1:2:(2*col)-1) = inter_hi;
[row,col] = size(intl);

X1 = conv2(intl,L0);
X2 = conv2(inth,HI);

% add the two together

BACKa = X1 + X2;

l = length(L0);
orig_row = 2*nrow - 1;
orig_col = 2*ncol - 1;

BACK = BACKa(l:orig_row + 1 - 1,l:orig_col + 1 - 1)*.25;

```


Appendix B. Microcalcification Detection Code

B.1 Introduction

After the wavelet filtered image has been created, the next step is to use it to detect the actual microcalcification clusters. This appendix describes the set of subroutines that make up the actual detection portion of the procedure. The hi-level program `detect.m` controls the whole process from front to back, while calling several subroutines to make important calculations. The primary subroutines are `main_seg_roc.m`, `pre_proc64.m`, `hypoth.m`, `roi_correl64.m`, `reduce64.m`, `roc_feat64.m`, and `sorter64.m`. In addition, there are several minor subroutines that perform various odd jobs such as reading in files, computing centers of mass, etc., that are not listed.

B.2 main_seg_roc.m

This m-file segments the mammogram by generating the row and column coordinates of the centers of the initial ROIs. Inputs are the filtered image and the multiple of the standard deviation that is to be used for the threshold, and the outputs are the center row and column coordinates (in vector form) and the filtered image standard deviation, which is used for future reference.

```
function [I_clear,J_clear,sd] = main_seg_roc(IMG,mult);

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% main_seg.m
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% Program main_seg.m that executes the segmentation functions for a
% 1020 by 2028 reconstructed wavelet filtered and thresholded mammogram.

% The requirements to run this program are:
%
% 1: A 1020x2028 matrix called IMG exists in memory

% The program parameters are:
%
% top_margin uncertainty edge distance from top/bottom of IMG
% size_margin uncertainty edge distance from sides of IMG
% min_energy minimum "energy" required for ROI to be
% considered relevant after first pass
% thresh minimum energy to survive the second pass
```

```

% box_row # rows in the sliding window (size in rows)
% box_col # cols in the sliding window (size in cols)
% -NOTE: the (image size - margin) / box size
% must be an integer !!!

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Initial Threshold
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

sd = std(IMG(:))
hi_t = mult*sd;
MASK = IMG > hi_t;
IMG = IMG.*MASK;
clear MASK

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Parameter Definitions
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

top_margin = 30;
side_margin = 22;
box_row = 64;
box_col = 64;

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% BEGIN PROGRAM
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% Compute the "Energy" matrix E

E = slider(IMG,top_margin,side_margin,box_row,box_col);

% Keep only those regions which have at least the minimum energy

min_energy = 4*hi_t;
[I,J] = find(E > min_energy);
I_mid = (I-1)*box_row+top_margin+(box_row/2);
J_mid = (J-1)*box_col+side_margin+(box_col/2);

% Perform the centroid migration

[G,EN] = SEG(IMG,I,J,top_margin,side_margin,min_energy,box_row,box_col);

thresh = 3*min_energy;
[I_final,J_final,E_final] = reducer(G,EN,thresh);

for i = 1:length(I_final)
if (I_final(i) < (box_row/2) | I_final(i) > (1020 - (box_row/2)))
E_final(i) = 0;
elseif (J_final(i) < (box_col/2) | J_final(i) > (2028 - (box_col/2)))

```

```

E_final(i) = 0;
end
end

F = find(E_final);
for i = 1:length(F);
I_clear(i) = I_final(F(i));
J_clear(i) = J_final(F(i));
E_clear(i) = E_final(F(i));
end

```

B.3 pre_proc64.m

This m-file implements the low and hi thresholding and identifies the individual objects within a ROI. The input is the filtered ROI, and the output is the ROI with only the separate objects that are potential microcalcifications remaining. The subroutine `roi_cut64.m` actually designates individual objects within the ROI, so it included here.

```

function ROIB = pre_proc64(ROI,tl,th,min_size);

%*****
%
% function ROIB = pre_proc64(ROI,tl,th,min_size);
%
% m-file returns the matrix ROIB which contains only those
% pixels which are part of an object that has at least one
% pixel with a value >= th.
%
%*****

% create masks of objects that have all pixels in excess of
% a given value (th or tl), and are at least min_size large

MASK1 = roi_cut64(ROI,th,min_size);
MASK2 = roi_cut64(ROI,tl,min_size);
Y = MASK1;

e_new = sum(MASK1(:));
dif = e_new;

% grow the ROI by dilating and ANDing the two masks

while dif > 0,
e_old = e_new;
X = Y>0;
BW = dilate(X,'fatten');

```

```

Y = BW.*MASK2;
e_new = sum(Y(:));
dif = e_new - e_old;
end

```

```

ROIB = Y;

```

```

function BACK = roi_cut(IMG,mult,min_size);

```

```

%*****
%
% function BACK = roi_cut(IMG,mult,min_size);
%
% Function that returns ROI BACK with only objects that have all
% of their pixel values in excess of mult and are at least min_size
% pixels large
%
%*****

```

```

MASK = IMG > mult;
X = MASK.*IMG;

```

```

[I,J,V] = find(abs(X));
l = length(I);

```

```

% first pass - assign same row clusters

```

```

C(1) = 1;
cmax = 1;
cind = cmax;
for i = 2:l
    new_col = J(i) - J(i-1);
    if new_col == 0
        t = find((I == I(i)) & (J == (J(i) - 1)));
        if t == []
            if I(i) == (I(i-1) + 1)
                C(i) = cind;
            else
                cmax = cmax + 1;
                cind = cmax;
                C(i) = cind;
            end
        else
            cind = C(t);
            C(i) = cind;
        end
    elseif new_col == 1
        t = find((I == I(i)) & (J == (J(i) - 1)));
        if t == []
            cmax = cmax + 1;

```

```

cind = cmax;
C(i) = cind;
else
cind = C(t);
C(i) = cind;
end
else
cmax = cmax + 1;
cind = cmax;
C(i) = cind;
end
end

```

```

% second pass - assign same column clusters

```

```

for i = 2:1
if (J(i) == J(i-1)) & (I(i) == I(i-1)+1)
if C(i) ~= C(i-1)
t = C(i-1);
T = find(C == t);
q = length(T);
for k = 1:q
C(T(k)) = C(i);
end
end
end
end

```

```

% determine the number of unique clusters, size, energy & distance

```

```

num = 0;
for i = 1:cmax
T = find(C == i);
if T ~= []
num = num + 1;
s = length(T);
if s > (min_size-1)
RI = [RI; I(T)];
RJ = [RJ; J(T)];
RE = [RE; V(T)];
end
end
end

```

```

SP = sparse(RI,RJ,RE,64,64);
BACK = full(SP);

```

B.4 hypoth.m

This m-file computes the features of size and linearity of objects within a ROI that are used as a first cut to separate microcalcifications from noise.

```
function [sz,LIN] = hypoth(ROIX);

%*****
%
% function [sz,LIN] = hypoth(ROIX);
%
% Returns the size and linearity for each object within ROIX
%
%*****

% Group the pixels into their corresponding objects - see roi_cut64.m

[I,J,C] = sep_cluster(ROIX);

set = -pi/2:pi/50:pi/2;
cosx = cos(set);
sinx = sin(set);

cmax = max(C);
for k = 1:cmax
    T=find(C==k);
    if (T~=[] & I ~= [])
        sz(k) = length(T);
        II = I(T);
        JJ = J(T);

        % *** Evaluate 'linearness' of the potential uCa++

        p1 = II*sinx;
        p2 = JJ*cosx;
        huff = p1 + p2;
        std_huff = std(huff);
        vhuff = min(std_huff);
        vhuff =vhuff^2;
        LIN(k) = vhuff/length(II);

    else
        sz(k) = 0;
        LIN(k) = 0;
    end
end
```

B.5 roi_correl64.m

This m-file generates the correlation peak and mean-square error values for each object within a ROI. The original and filtered ROIs are passed as input, with vectors of the error values for each object being created as output.

```
function [middle,msq] = roi_correl64(X,PIC);

%*****
%
% function [RX,AC] = roi_correl64(X,PIC);
%
% wavelet roi = X (pre-processed)
% original roi = PIC
% This m-file evaluates every cluster in X and returns its central correlation
% value (middle), and its correlation mean-square error (msq)
%
%*****

PIC2 = PIC - mean(PIC(:));
NPIC = PIC2>0;
NPIC = NPIC.*PIC2;

[I,J,C] = sep_cluster(X);
cmax = max(C);
for k = 1:cmax;
T = find(C == k);
if (T ~= [] & I ~= [])
II = I(T);
JJ = J(T);
min_row = min(II);
min_col = min(JJ);
max_row = max(II);
max_col = max(JJ);
MI = II - min_row + 1;
MJ = JJ - min_col + 1;
l = length(II);
Ms = ones(l,l);
S = sparse(MI,MJ,Ms,(max_row-min_row+1),(max_col-min_col+1));
TMP = sparse(II,JJ,Ms,64,64);
TTMP = full(TMP);
MASK = full(S);
MASK = MASK/sqrt(sum(MASK(:)));
AC = xcorr2(MASK);
[nr,nc] = size(AC);
midr = ceil(nr/2);
midc = ceil(nc/2);
top = midr - 3;
bot = midr + 3;
```

```

if top < 1
top = 1;
bot = nr;
end
left = midc - 3;
right = midc + 3;
if left < 1
left = 1;
right = nc;
end
midac = AC(top:bot,left:right);
if min_row < 5
top = 1;
toffset = -(5 - min_row);
else
top = min_row - 4;
toffset = 0;
end
if max_row > 60
bot = 64;
toffset = max_row - 60;
else
bot = max_row + 4;
toffset = 0;
end
if min_col < 5
left = 1;
soffset = -(5 - min_col);
else
left = min_col - 4;
soffset = 0;
end
if max_col > 60
right = 64;
soffset = max_col - 60;
else
right = max_col + 4;
soffset = 0;
end
piece = NPIC(top:bot,left:right);
if toffset < 0
piece = [zeros(abs(toffset),right-left+1); piece];
elseif toffset > 0
piece = [piece; zeros(toffset,right-left+1)];
end
[nr,nc] = size(piece);
if soffset < 0
piece = [zeros(nr,abs(soffset)) piece];
elseif soffset > 0
piece = [piece zeros(nr,soffset)];
end

```



```

pp = sqrt(sum(sum(piece.^2)));
if pp > 0
piece = piece/pp;
else
piece = piece;
end
[nr,nc] = size(piece);
midr = ceil(nr/2) + toffset;
midc = ceil(nc/2) + soffset;
CC = xcorr2(piece,MASK);
if max(CC(:)) > 0
CC = CC/max(CC(:));
else
CC = CC;
end
% extract middle section of each correlation
[nr,nc] = size(CC);
midr = ceil(nr/2);
midc = ceil(nc/2);
[nr,nc] = size(midac); % should be 1,3,5, or 7
rmar = (nr-1)/2; % should be 0,1,2, or 3
cmar = (nc-1)/2;
midcc = CC((midr-rmar):(midr+rmar),(midc-cmar):(midc+cmar));
% compare how similar the vectors are
xhat = max(midcc(:));
mm = min(midcc(:));
if xhat > mm
yy = xhat/(xhat - mm);
else
yy = 0;
end
newcc = yy*(midcc - mm);
[nr,nc] = size(newcc);
midr = ceil(nr/2);
midc = ceil(nc/2);
middle(k) = newcc(midr,midc);
% compute diff based on 5x5 piece of correlation values
[nr,nc] = size(newcc);
if (nr > 5 & nc > 5)
newcc = newcc(2:6,2:6);
midac = midac(2:6,2:6);
end
diff = newcc - midac;
msq(k) = mean(mean(diff.^2));
else
middle(k) = 0;
msq(k) = 1;
end
end
end

```

B.6 reduce64.m

This m-file removes the non-microcalcifications from an ROI by retaining only those objects whose index number is passed in the vector 'keepers'.

```
function NEW = reduce64(X,keepers);

%*****
%
% function NEW = reduce64(X,keepers);
%
% Returns the ROI with only those objects indexed by the vector
% keepers
%
%*****

[I,J,C] = sep_cluster(X);
l=length(keepers);
for k=1:l
    T = find(C == keepers(k));
    II = I(T);
    JJ = J(T);
    IT = [IT; II];
    JT = [JT; JJ];
end
lo=length(IT);
Ms=ones(lo,1);
S=sparse(IT,JT,Ms,64,64);
TMP=full(S);
NEW=TMP.*X;
```

B.7 roc_feat64.m

This m-file cross compares the reduced wavelet filtered ROI with the hi-pass filtered ROI and generates the area and area/perimeter features.

```
function feat = roc_feat64(NEW,M,lt);

%*****
%
% function feat = roc_feat64(NEW,M,lt);
%
% function compares hi-pass and wavelet filtered ROIs and outputs
% size and area/perim features for each microcalcification
%
```

```
%*****
```

```
MS = NEW>0;
np = .9*sum(sum(MS));

ut = .9*max(M(:));
G=pre_proc64b(M,lt,ut,4);
MK = G>0;
tot = sum(sum(MK));

while (tot < np & ut > lt),
    ut = ut - 5;
    G=pre_proc64b(M,lt,ut,4);
    MK = G>0;
    tot = sum(sum(MK));
end

[sz,lin]=hypoth(G);
T=find(sz>5 & lin > .025);
PATCH = reduce64(G,T);

MK = PATCH>0;
DIF = MS.*MK;

BW = dilate(DIF,'dilate');
XXX = BW.*PATCH;
e_new = sum(XXX(:));
dif = e_new;

while dif > 0,
    e_old = e_new;
    Y = XXX > 0;
    BW = dilate(Y,'dilate');
    XXX = BW.*PATCH;
    e_new = sum(sum(XXX));
    dif = e_new - e_old;
end

VEC = [];
[I,J,C] = sep_cluster(XXX);
cmax = max(C);
cnt = 0;
for k = 1:cmax;
    T = find(C == k);
    if (T ~= [] & I ~=[])
        cnt = cnt + 1;
        SIZE = length(T);
        II = I(T);
        JJ = J(T);
        l = length(II);
        Ms = ones(l,1);
```

```

TMP = sparse(II,JJ,Ms,64,64);
MASK = full(TMP);
perim = bwperim(MASK);
PERIM = sum(sum(perim));
COMPACT = SIZE/PERIM;
VEC = [VEC; cnt SIZE COMPACT];
end
end

feat=VEC;

```

B.8 sorter64.m

This m-file examines the list of features produced by roc_feat64.m to determine which ROIs contain 3 or more valid microcalcifications.

```

function F = sorter64(FEAT);

%*****
%
% function F = sorter64(FEAT);
%
% function that evaluates the list of features FEAT to determine
% which ROIs contain 3 or more valid microcalcifications, and
% outputs F as a list of valid ROIs
%
%*****

mx = max(FEAT(:,1));
KEEP = [];
F=[];

for k = 1:mx
    T = find(FEAT(:,1) == k);
    if T ~= []
        SUB = FEAT(T,:);
        nr = length(T);
        kp = 0;
        for n = 1:nr
            sz = SUB(n,3);
            rat = round(sz/SUB(n,4));
            inter = sz - rat;
            if sz > 5 & sz < 18 & inter > 0
                kp = kp + 1;
            elseif sz > 17 & sz < 22 & inter > 1
                kp = kp + 1;
            elseif sz > 21 & sz < 29 & inter > 2

```

```
kp = kp + 1;
elseif sz > 28 & sz < 35 & inter > 3
kp = kp + 1;
elseif sz > 34 & sz < 44 & inter > 4
kp = kp + 1;
elseif sz > 43 & sz < 46 & inter > 5
kp = kp + 1;
end
end
if kp > 2
KEEP = [KEEP; k];
F = [F; SUB];
end
end
end
```

Appendix C. Feature Selection Code

C.1 Introduction

This appendix contains the code used to calculate the Effective Decision Boundary Feature Matrix Σ_{EDBFM} . Four separate m-files are used to compute this matrix: `classified.m` prunes the input data so that only those data points are retained which are on the correct side of the decision boundary. It calls `parzen_prob_fast.m` for each point to compute a nonparametric probability estimate for each point using a Parzen Window pdf estimate. Once the data is correctly classified, `find_bound.m` computes the decision boundary points. Then `get_edbfm.m` is executed to compute the matrix Σ_{EDBFM} . The dominant features for classification purposes can then be found by computing the eigenvalues of Σ_{EDBFM} .

C.2 `classified.m`

```
function Qx_prime = classified(Qx,Qy,h);

% function Qx_prime = classified(Qx,Qy,h)
%
% function that retains only those members of class Qx in Qx_prime that
% are correctly classified (based on a Bayes decision rule) using a
% Parzen window pdf estimate.
% h is the parzen window width parameter

icvx = inv(cov(Qx));
icvy = inv(cov(Qy));
sigx = sqrt(det(cov(Qx)));
sigy = sqrt(det(cov(Qy)));
l = size(Qx,1);           % # of rows

for i = 1:l
    f = Qx(i,:);
    px(i) = parzen_prob_fast(Qx,f,h,icvx,sigx);
    py(i) = parzen_prob_fast(Qy,f,h,icvy,sigy);
end
dist = px./py;
dist = dist(:);
keep = find(dist>1);
s = size(keep,1);         % # of retained elements
for i = 1:s
    p = keep(i);
    Qx_prime(i,:) = Qx(p,:);
end
```

C.3 parzen_prob_fast.m

```
function prob = parzen_prob_fast(Qx,f,h,icvx,sig);

% function prob = parzen_prob_fast(Qx,f,h,icvx,sig);
%
% function that computes the a posteriori probability estimate that
% vector f is a member of the class defined by the data set Qx. h is
% the parzen window width parameter. The data Qx should be "tall and
% skinny" - i.e., every column is a feature and every row a realization
% of the data set. f should be a row vector
% the inverse cov matrix (icvx) and sqrt of the determinant of the cov matrix
% (sig) are also passed so they don't have to be computed too much

n = size(Qx,2);      % # columns - dimensionality of the data
s = size(Qx,1);      % # rows - # data points (members) of class Qx

wun = ones(s,1);
big_f = wun*f;        % replicate f s times to form the big matrix

X = big_f - Qx;        % x - x_hat for the mahal calculation

dist = diag((X*icvx*X')/(2*h^2));

dist = sum(exp(-dist));
prob = dist/(s*(2*pi)^(n/2)*sig*h^(n+1));
```

C.4 find_bound.m

```
function [B0,Yj] = find_bound(Qx,Qy,h,Lmin,thresh);

% function B0 = find_edbfm(Qx,Qy,h,Lmin,thresh);
%
% function that takes in the two classes Qx and Qy and finds the edbfm
% matrix and the decision boundary B0. h is the parzen window width
% parameter and Lmin is the number of points in class 2 closest to class 1
% that are used to compute the decision boundary (usually ~ 5)
% thresh is the decision boundary 'nearness' threshold
%

% retain only the member of each class that are on the correct side
% of the decision boundary

Qx_ok = classified(Qx,Qy,h);
Qy_ok = classified(Qy,Qx,h);

% pick the Lmin members of class 2 that are closest (probabilistic) to class 1
```

```

l = size(Qy_ok,1);
icvx = inv(cov(Qx));
icvy = inv(cov(Qy));
sigx = sqrt(det(cov(Qx)));
sigy = sqrt(det(cov(Qy)));
for i = 1:l
    y = Qy_ok(i,:);
    d(i) = parzen_prob_fast(Qx,y,h,icvx,sigx);
end

[yy,ii] = sort(d);
ii = ii(:);
ii = flipud(ii);      % resort so distances are in descending, rather than
                      % ascending order
Yj(1:Lmin,:) = Qy_ok(ii(1:Lmin),:);

% cycle thru each point in class 1 and find the closest point in the Lmin
% set of class 2, find the decision boundary point, and then compute
% the gradient

clear d
l = size(Qx_ok,1);

for i = 1:l
    i
    x = Qx_ok(i,:);
    for k = 1:Lmin
        near = Yj(k,:);
        d(k) = norm(x - near);
    end
    [yy,ii] = sort(d);
    closest = Yj(ii(1),:);
    anchor1 = x;
    anchor2 = closest;
    m = (anchor1 + anchor2)/2;
    t1 = parzen_prob_fast(Qx,m,h,icvx,sigx);
    t2 = parzen_prob_fast(Qy,m,h,icvy,sigy);
    t = -log(t1/t2);
    while abs(t) > thresh,
        if t > 0
            anchor2 = m;
            anchor1 = anchor1;
        else
            anchor1 = m;
            anchor2 = anchor2;
        end
        m = (anchor1 + anchor2)/2;
        t1 = parzen_prob_fast(Qx,m,h,icvx,sigx);
        t2 = parzen_prob_fast(Qy,m,h,icvy,sigy);
        t = -log(t1/t2);
    end
end

```



```

        B0(i,:) = m;
end

```

C.5 get_edbfm.m

```

function EDBFM = get_edbfm(Qx,Qy,B0,h,delta);

% function EDBFM = get_edbfm(Qx,Qy,B0,h,delta);
%
% computes the EDBFM for data of classes Qx and Qy, with a decision
% boundary B0. It uses a parzen window pdf estimate with width parameter
% h and gradient increment of delta to numerically find the normal
% vector

icvx = inv(cov(Qx));
icvy = inv(cov(Qy));
sigx = sqrt(det(cov(Qx)));
sigy = sqrt(det(cov(Qy)));
m = size(B0,1)      % # rows
n = size(B0,2)      % # columns

N = zeros(n,n);
for i = 1:m
    X = B0(i,:);
    t1 = parzen_prob_fast(Qx,X,h,icvx,sigx);
    t2 = parzen_prob_fast(Qy,X,h,icvy,sigy);
    tX = -log(t1/t2);
    for k = 1:n
        f = X;
        f(k) = f(k) + delta;
        t1 = parzen_prob_fast(Qx,f,h,icvx,sigx);
        t2 = parzen_prob_fast(Qy,f,h,icvy,sigy);
        t = -log(t1/t2);
        grad_h(k) = (tX - t)/delta;
    end
    normal = grad_h'*grad_h;
    N = N + normal;
end
EDBFM = N;

```

Bibliography

1. *Breast Cancer: New Perspectives Can Replace Unrealistic Fears*. Technical Report ISSN 0741-6254, Rochester, MN: Mayo Foundation for Medical Education and Research, October 1994.
2. *Model-Driven Automatic Target Recognition*. Technical Report, Wright-Patterson AFB OH: ARPA, October 1994.
3. Brodie, I. and R. A. Gutcheck. "Radiographic Information Theory and Application to Mammography," *Medical Physics*, 9 (1982).
4. Chang, Tianhorng and C.C. Jay Kuo. "Texture Analysis and Classification with Tree-Structured Wavelet Transform," *IEEE Transactions on Image Processing*, 2(4) (October 1993).
5. Cohen, A., et al. *Biorthogonal Bases of Compactly Supported Wavelets*. Technical Report TM 11217-900529-07, AT&T Bell Lab, 1990.
6. Daubechies, I. "Ten Lectures on Wavelets," *CBMS-NSF Regional Conference Series in Applied Mathematics*, (61) (1992).
7. Dhawan, Atam P. "Computerized Mammographic Image Analysis for Reducing False Positive Rate for Biopsy Recommendation," *SPIE*, 1905:540-541 (1993).
8. Dhawan, Atam P, et al. "Artificial Neural Network Based Classification of Mammographic Microcalcifications Using Image Structure Features," *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1377-1401 (1993).
9. Dijkerman, Robert W. and Ravi R. Mazumdar. "Wavelet Representations of Stochastic Processes and Multi-resolution Stochastic Models," *IEEE Transactions on Signal Processing*, 42(7) (July 1994).
10. Duda, Richard O. and Peter E. Hart. *Pattern Classification and Scene Analysis*. New York: John Wiley & Sons, 1973.
11. Eisenbies, Christopher. *Classification of Ultra High Resolution Radar Using Decision Boundary Analysis*. MS thesis, AFTT/EN/ENG/94D-04, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, 1994.
12. Giger, Maryellen. "Computer-Aided Diagnosis," *RSNA Categorical Course in Physics*, 283-298 (1993).
13. Giger, Maryellen and Carl J. Vyborny. "Computer Vision and Artificial Intelligence in Mammography," *AJR*, 162:699-708 (March 1994).
14. Gonzalez, Rafael C. and Paul Wintz. *Digital Image Processing*. Reading, Mass: Addison Wesley, 1987.
15. Gordon, Richard and Rangaraj M. Rangayyan. "Feature Enhancement of Film Mammograms Using Fixed and Adaptive Neighborhoods," *Applied Optics*, 23(4):560-564 (February 1984).
16. Jawerth, Bjorn D., et al. "Compressing Still and Moving Images with Wavelets," *Multimedia Systems*, 2(3) (1994).
17. Karssemeijer, N. "Recognition of Clustered Microcalcifications Using a Random Field Model," *Proceedings of the SPIE*, 1905:776-786 (1993).

18. Kocur, Catherine M. *Computer-Aided Breast Cancer Diagnosis*. MS thesis, AFIT/GSO/ENS/94-D-03, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, 1994.
19. Laine, Andrew, et al. "Wavelets for Contrast Enhancement of Digital Mammography," *IEEE Engineering in Medicine and Biology Magazine*, 14(5) (September 1995).
20. Laine, Andrew and Sergio Schuler. "Hexagonal Wavelet Processing of Digital Mammography," *SPIE Image Processing*, 1898 (1993).
21. Laine, Andrew, et al. "Mammographic Feature Enhancement by Multiscale Analysis," *IEEE Transactions on Medical Imaging*, 13(4) (December 1994).
22. Lee, Chulhee and David Langrebe. "Decision Boundary Feature Extraction for Nonparametric Classification," *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2) (April 1993).
23. Lee, Chulhee and David Langrebe. "Feature Extraction Based on Decision Boundaries," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(4) (April 1993).
24. Mallat, S. G. "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7) (1989).
25. Morrow, W. M., et al. "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, 11(3) (September 1992).
26. Myers, Lemuel. *Image Interpretation and Enhancement for the Visually Impaired*. MS thesis, AFIT/GEO/ENG/94D-55, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, 1994.
27. Qian, Wei, et al. "Computer Assisted Diagnosis for Digital Mammography," *IEEE Engineering in Medicine and Biology Magazine*, 14(5) (September 1995).
28. Quan, Wei, et al. "Digital Mammography: M-Channel Quadrature Mirror Filters for Microcalcification Extraction," *Computerized Medical Imaging and Graphics*, 18(5) (October 1994).
29. Quan, Wei, et al. "Tree-Structured Nonlinear Filters in Digital Mammography," *IEEE Transactions on Medical Imaging*, 13(1) (March 1994).
30. Rioul, Olivier and Martin Vetterli. "Wavelets and Signal Processing," *IEEE SP Magazine*, 1053-5888/91:14-38 (October 1991).
31. Smiley, Steven. E. *Image Segmentation Using Affine Wavelets*. MS thesis, AFIT/EN/ENG/91D-50, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, 1991.
32. Smith, R. A. "Epidemiology of Breast Cancer," *RSNA Categorical Course in Physics*, 21-33 (1993).
33. Stewart, John. *Stock Market Prediction*. MS thesis, AFIT/EN/ENG/94D-04, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, 1995.
34. Strickland, R. N. and H. I. Hahn. "Detection of Microcalcifications using Wavelets," *Digital Mammography: Proceedings of the 2nd International Workshop on Digital Mammography* (July 1994).
35. Tanne, Janice Hopkins. "Everything You Need to Know About Breast Cancer...But Were Afraid to Ask," *New York [GNYC]*, 26:52-62 (October 1993).

36. Vaidyanathan, P. P. and Phuong-Quan Hoang. "Lattice Structures for Optimal Design and Robust Implementation of Two-Channel Perfect Reconstruction QMF Banks," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(1) (January 1988).
37. Yaffe, Martin J. "Digital Mammography," *RSNA Categorical Course in Physics*, 271–282 (1993).
38. Yoshida, Hiroyuki, et al. "Automated Detection of Clustered Microcalcifications in Digital Mammograms Using Wavelet Transform Techniques," *SPIE Image Processing*, 2167:868–886 (1994).

Vita

Lieutenant Donald A McCandless was born and raised in the Seattle, Washington area. In 1991 he received his Bachelor of Science in Electrical Engineering from the University of Washington and was commissioned a 2nd Lieutenant in the United States Air Force. Prior to coming on active duty he worked for the U.S. Army at Dugway Proving Ground in western Utah as an instrumentation engineer for chemical, smoke, and illumination tests. Called to active duty in September 1992, he served at the National Air Intelligence Center as an Air Defense Command, Control, and Communications analyst until May 1994, when he entered AFIT as a full time student specializing in communications and signal processing. Upon completion of his Master of Science degree, Lt. McCandless will move with his wife and son to Albuquerque, New Mexico where he will work at the Air Force's Phillips Lab.

Permanent address: 44 Wells Dr
Dayton, OH 45431

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 1995		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE DETECTION OF CLUSTERED MICROCALCIFICATIONS USING WAVELETS			5. FUNDING NUMBERS	
6. AUTHOR(S) Donald A. McCandless 1st Lieutenant, USAF				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, WPAFB OH 45433-6583			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GE/ENG/95D-17	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) NA			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; Distribution Unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) An automated method for detecting microcalcification clusters is presented. The algorithm begins with a digitized mammogram and outputs the center coordinates of regions of interest (ROIs) that contain microcalcification clusters. The method presented uses a 12-tap Least Asymmetric Daubechies (LAD12) wavelet in a tree structured filter bank to increase the signal to noise level of microcalcifications. The signal to noise level gain achieved by the filtering allows subsequent thresholding to eliminate on average 90% of the image from further consideration without eliminating actual microcalcifications 95% of the time. A novel approach to isolating individual calcifications from background tissue through non-stationary noise reduction, low/hi thresholding, and morphological filtering is demonstrated; this technique reduces the number of false detections by an average of 5 per image. Several features are extracted from each potential calcification, including two newly proposed correlation features, to distinguish actual microcalcifications from correlated background tissue. Altogether, the method successfully detected 44 of 53 microcalcification clusters (83%) with an average of 2.3 false positive clusters per image. A cluster is considered detected if it contains 3 or more microcalcifications within a 6.4 mm by 6.4 mm area. Although the emphasis is placed on detecting microcalcification clusters for further examination by a radiologist with no attempt made to diagnose the cluster as malignant or benign, the method successfully detected 13 of 14 (93%) malignant cases.				
14. SUBJECT TERMS Wavelets, Pattern Recognition, Breast Cancer, Medical Imaging, Model Based Vision, Microcalcifications			15. NUMBER OF PAGES 78	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to **stay within the lines** to meet **optical scanning requirements**.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."

DOE - See authorities.

NASA - See Handbook NHB 2200.2.

NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.

DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.

NASA - Leave blank.

NTIS - Leave blank.

Block 13. Abstract. Include a brief (*Maximum 200 words*) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (*NTIS only*).

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.